

CONSCIOUSNESS

AN INTRODUCTION

THIRD EDITION

SUSAN BLACKMORE
EMILY T. TROSCIANKO



Consciousness

Is there a theory that explains the essence of consciousness?

Or is consciousness itself an illusion?

Am I conscious now?

Now considered the 'last great mystery of science', consciousness was once viewed with extreme scepticism and was rejected by mainstream scientists. It is now a significant area of research, albeit a contentious one, as well as a rapidly expanding area of study for students of psychology, philosophy, and neuroscience.

This edition of *Consciousness*, revised by author team Susan Blackmore and Emily Troscianko, explores the key theories and evidence in consciousness studies ranging from neuroscience and psychology to quantum theories and philosophy. It examines why the term 'consciousness' has no recognised definition and provides an opportunity to delve into personal intuitions about the self, mind, and consciousness.

Featuring comprehensive coverage of all core topics in the field, this edition includes:

- Why the problem of consciousness is so hard
- Neuroscience and the neural correlates of consciousness
- Why we might be mistaken about our own minds
- The apparent difference between conscious and unconscious
- Theories of attention, free will, and self and other
- The evolution of consciousness in animals and machines
- Altered states from meditation to drugs and dreaming

Complete with key concept boxes, profiles of well-known thinkers, and questions and activities suitable for both independent study and group work, *Consciousness* provides a complete introduction to this fascinating field. Additional resources are available on the companion website: www.routledge.com/cw/blackmore

Susan Blackmore is a psychologist, TED lecturer, and writer researching consciousness, memes, meditation, and anomalous experiences, and is Visiting Professor in Psychology at the University of Plymouth. *The Meme Machine* (1999) has been translated into 16 languages; more recent books include *Zen and the Art of Consciousness* (2011) and *Seeing Myself: The New Science of Out-of-Body Experiences* (2017).

Emily T. Troscianko is a writer and researcher interested in mental health, readers' responses to literature, and how the two might be linked – as well as what both have to do with human consciousness. She is a Research Associate at The Oxford Research Centre in the Humanities (TORCH), University of Oxford, writes the blog 'A Hunger Artist' for *Psychology Today*, and has published a monograph, *Kafka's Cognitive Realism* (2014), exploring the strange phenomenon we call the 'Kafkaesque'.

Consciousness

An Introduction

3rd Edition

Susan Blackmore and
Emily T. Troscianko

First published 2018
by Routledge
2 Park Square, Milton Park, Abingdon, Oxon OX14 4RN

Routledge is an imprint of the Taylor & Francis Group, an informa business

© 2018 Susan Blackmore and Emily T. Troscianko

The right of Susan Blackmore and Emily T. Troscianko to be identified as authors of this work has been asserted by them in accordance with sections 77 and 78 of the Copyright, Designs and Patents Act 1988.

All rights reserved. No part of this book may be reprinted or reproduced or utilised in any form or by any electronic, mechanical, or other means, now known or hereafter invented, including photocopying and recording, or in any information storage or retrieval system, without permission in writing from the publishers.

Trademark notice: Product or corporate names may be trademarks or registered trademarks, and are used only for identification and explanation without intent to infringe.

British Library Cataloguing-in-Publication Data

A catalogue record for this book is available from the British Library

All efforts have been made to secure permissions for items belonging to third parties. The Publisher would be glad to hear of any instance where items have not been duly acknowledged.

First edition published 2004 by Oxford University Press

ISBN: 978-1-138-80130-1 (hbk)


ISBN: 978-1-138-80131-8 (pbk)

ISBN: 978-1-315-75502-1 (ebk)

Typeset in Myriad Pro
by Apex CoVantage, LLC

Visit the companion website: www.routledge.com/cw/blackmore

*To all the students who took Sue's
consciousness course.*




Contents

Boxes	ix
Prefaces	xi
Acknowledgements	xv
Introduction	1
Section One The problem	9
1 What's the problem?	11
2 What is it like to be . . . ?	32
3 The grand illusion	52
Section Two The brain	75
4 Neuroscience and the correlates of consciousness	77
5 The theatre of the mind	103
6 The unity of consciousness	128
Section Three Body and world	157
7 Attention	159
8 Conscious and unconscious	186
9 Agency and free will	218
Section Four Evolution	247
10 Evolution and animal minds	249
11 The function of consciousness	276
12 The evolution of machines	303
Section Five Borderlands	341
13 Altered states of consciousness	343
14 Reality and imagination	372
15 Dreaming and beyond	398

• C O N T E N T S

Section Six Self and other	433
16 Egos, bundles, and theories of self	435
17 The view from within?	464
18 Waking up	490
References	513
Index	599




Boxes

Chapter	Profile	Concept	Practice	Activity
Introduction	0.1 Blackmore 2 0.2 Troscianko 4			
1 What's the problem?	1.1 Descartes 16 1.2 James 22	1.1 The hard problem 29	1.1 Am I conscious now? 13	1.1 Defining consciousness 21
2 What is it like to be ...?	2.1 Chalmers 41 2.2 Churchland 49	2.1 The philosopher's zombie 44	2.1 What is it like to be me now? 35	2.1 Mary the colour scientist 39
3 The grand illusion	3.1 Ramachandran 60	3.1 Magic 67 3.2 Seeing or blind? 70	3.1 How much am I seeing now? 55	3.1 Filling-in 62
4 Neuroscience and the neural correlates of consciousness	4.1 Koch 95	4.1 Mapping the brain 79 4.2 Phantom phenomena 98	4.1 Where is this pain? 97	4.1 The rubber hand illusion 101
5 The theatre of the mind	5.1 Dennett 104 5.2 Baars 113	5.1 Seeing blue 111	5.1 What is it that is conscious? 106	5.1 Cartesian materialism 124
6 The unity of consciousness	6.1 Tononi 138	6.1 Synaesthesia 143 6.2 Orwellian and Stalinesque 155	6.1 Is this experience unified? 130	6.1 Are you a synaesthete? 145 6.2 Split-brain twins 149 6.3 The cutaneous rabbit 154
7 Attention	7.1 Graziano 168		7.1 Did I direct my attention? 161	7.1 Meditation 177
8 Conscious and unconscious	8.1 Goodale 203 8.2 Clark 214	8.1 Sensory substitution 208	8.1 Did I do this consciously? 189 8.2 Was this decision conscious? 211	8.1 Incubation 201

• B O X E S

Chapter	Profile	Concept	Practice	Activity
9 Agency and free will	9.1 Wegner 238	9.1 Volition and timing 232	9.1 Am I doing this? 218	9.1 Getting out of bed 222 9.2 Libet's voluntary action 229
10 Evolution and animal minds	10.1 Dawkins 255 10.2 Grandin 262	10.1 Deception 267	10.1 What is it like to be that animal? 258	10.1 Lab choice 273
11 The function of consciousness	11.1 Humphrey 288	11.1 Four ways of thinking about the evolution of consciousness 284 11.2 Memes 298	11.1 Am I conscious now? Does this have a function? 281 11.2 Is this a meme? 300	11.1 The sentence line 286
12 The evolution of machines	12.1 Turing 307 12.2 Searle 326 12.3 Holland 332	12.1 Brains and computers compared 308 12.2 Humanoid robots and simulations 323	12.1 Am I a machine? 314 12.2 Is this machine conscious? 325	12.1 A Turing test for creativity 317 12.2 The Seventh Sally 321
13 Altered states of consciousness	13.1 Metzinger 350	13.1 State-specific sciences 353 13.2 Is hypnosis an ASC? 365	13.1 Is this my normal state? 345	13.1 Discussing ASCs 357
14 Reality and imagination	14.1 Siegel 380	14.1 The ganzfeld controversy 392	14.1 Living without psi 388	14.1 Telepathy tests 390
15 Dreaming and beyond	15.1 Hobson 402 15.2 Revonsuo 402	15.1 The evolution of dreaming 403 15.2 Sleep paralysis 413	15.1 Staying awake while falling asleep 411 15.2 Becoming lucid 417 15.3 What survives? 426	15.1 Discussing hypnagogia 412 15.2 Inducing lucid dreams 416
16 Egos, bundles, and theories of self	16.1 Hume 438	16.1 Ego and bundle theories 439 16.2 Selves, clubs, and universities 446	16.1 Who is conscious now? 438 16.2 Am I the same 'me' as a moment ago? 448	16.1 The teletransporter 445
17 The view from within?	17.1 Varela 473	17.1 Do we need a new kind of science? 466	17.1 Is there more in my P-consciousness than I can access? 469 17.2 Solitude 482	17.1 Positioning the theories 477
18 Waking up	18.1 Harris 491	18.1 Koans 500 18.2 Pure consciousness 503	18.1 What is this? 495 18.2 Mindfulness 507	18.1 The headless way 499



Prefaces

PREFACE TO THE FIRST EDITION

I have loved writing this book. For many years, working as a lecturer, I never seemed to have enough time to read or think or do the work I really wanted to do. So in September 2000 I left my job and threw myself into the vast and ever-expanding literature of consciousness studies. Writing the book meant spending over two years mostly at home completely by myself, reading, thinking, and writing, which was a real pleasure.

I could never have worked this way without three things. First, there are all the conferences at which I have met other scientists and philosophers and been able to share ideas and arguments. Second, there is the internet and email, which make it possible to keep in touch with colleagues all over the world instantly without moving from my own desk. Third, there is the WWW, which has expanded beyond all recognition in the few years since I first thought of writing this book. I am constantly amazed at the generosity of so many people who give their time and effort to make their own work, and the work of others, freely available to us all.

I could never have enjoyed working at home so much were it not for my wonderful family: my partner Adam Hart-Davis and my two children Emily and Jolyon Troscianko. Having Joly drawing the cartoons meant many happy battles over whether self is more like a candle, a raindrop, or bladderwrack seaweed, and what the Cartesian Theatre would look like if it existed. My thanks go to them all.

PREFACE TO THE SECOND EDITION

So much has happened in the past seven or eight years of consciousness studies! So updating this book has been a real challenge. Although there have been new philosophical ideas and some theoretical developments, the real impetus for change has come from neuroscience. Questions that, even a few years ago, seemed beyond empirical reach are now routinely being addressed by experiments.

One example especially dear to my heart is the out-of-body experience. Traditionally rejected by experimental psychologists as an oddity, or even make-believe, OBEs seemed to evade any theoretical grip. Back in the 1980s, when I was researching these strange experiences, most scientists agreed that nothing

actually left the body but, beyond vague speculation, could offer no convincing alternative. In the first edition of this book I described hints that an area of the temporal lobe might be implicated; now, in the second edition, I can describe repeatable experiments inducing OBEs, both by brain stimulation and by virtual reality methods. Theory has gone forward in leaps and bounds, and we can now understand how OBEs arise through failures of the brain mechanisms involved in constructing and updating the body image. As so often happens, learning about how something fails can lead to new insights into how it normally functions – in this case, our sense of bodily self.

There have been other new developments in the understanding of self. Not only are more philosophers learning about neuroscience and bringing these two disciplines closer together, but research in another previously fringe area – meditation – has provided surprising insights. From brain scans of long-term meditators, we can see how attentional mechanisms change after long training and how possibly the claim that self drops out may be grounded in visible brain changes.

In more down-to-earth ways, developments in machine consciousness have provided new constraints on how brains must work. Software and robot engineers struggle to make their systems do tasks that humans find easy and in the process are discovering what kinds of internal models and what kinds of embodiment and interactions with the outside world are, and are not, needed. It seems that we, like machines, build up ways of understanding our worlds that are completely impenetrable to anyone else – and this may give us clues to the nature of subjectivity and the apparent privacy and ineffability of qualia. All these discoveries feed into the various theories and increasingly mean they can be tested.

Then there is the great hunt for the neural correlates of consciousness. Personally, I think this highly active and popular approach is doomed to failure: it depends on the idea that some neural processes are conscious while others are not, and I believe this is nonsense. But I'm in a tiny minority here. The important thing is that this work will inevitably reveal which approach is right. The rapid pace of change over these past few years suggests that we may soon find out and makes the prospect of the next few years very exciting indeed.

I have changed, too. Since the first edition, I have written a *Very Short Introduction to Consciousness*, which, unlike this textbook, was explicitly meant to include my own ideas about consciousness. I enjoyed being made to explain so clearly why I think consciousness is an illusion. I then interviewed twenty top scientists and philosophers for my book *Conversations on Consciousness* and learnt that when Kevin O'Regan was a tiny boy he already thought of himself as a machine; that Ned Block thinks that O'Regan and Dennett don't even appreciate phenomenality; that Dan Dennett goes out of his way, every now and then, to give himself a good dose of the zombic hunch just so that he can practise abandoning it; and that Christof Koch, having thought so much about consciousness, doesn't squish bugs anymore. Having accepted that conscious will is an illusion, Dan Wegner said he gained a sense of peace in his life. Yet by contrast, most of my conversationalists, when asked 'Do you have free will?', said they did, or if not that they lived their lives as though they did, which is not something I feel I can do anymore.

Consciousness is an exciting subject – perhaps the most exciting mystery we can delve into now that neuroscience is giving us so many new tools. I have no idea

whether I will ever be able to update this book again. Even after so few years the task was daunting, and in a few more years the areas that seem important may have shifted completely. But we shall have to wait and see. Meanwhile, I hope you will enjoy battling with the great mystery.

PREFACE TO THE THIRD EDITION

SUE

As soon as I was invited to write a third edition I knew that the whole structure of this book would have to change. Indeed, I knew this back in 2009 when embarking, with both trepidation and enthusiasm, on the second. By then neuroscience was really beginning to take off, but I did manage to squeeze everything into the old scheme. By 2016 this was no longer feasible; there was just too much exciting new research to introduce, so what could I do? I am a lone worker. I rarely collaborate with others, and I love to work at home in silence and solitude. And even if I'd wanted to find a collaborator, who and where could they be, and how would we work together on such a complex book?

I was with my daughter in Oxford one day, sharing this huge problem with her, when we both spoke at once: – 'You wouldn't consider. . .?' – 'I could do it!' We laughed, and so our new collaboration was begun. I say 'collaboration' but in reality, Emily has done almost all the massive amount of work involved in bringing our book up to date. I gave advice, read and edited what she had done, and wrote some small pieces myself, but mostly what is new is her work. Her interest in language added new dimensions to the overview of consciousness studies; her deep understanding of eating disorders brought her knowledge of psychotherapy to bear; and her background in literary studies led to our including literary quotations in every chapter. I would never have thought of this and have found some of these excerpts quite moving – as well as thought-provoking.

Working within the family might have proved traumatic but did not. My husband, Adam Hart-Davis, supported us throughout. Vast differences in our academic backgrounds might have been a hindrance but instead seemed to be a help, and despite coming at the study of consciousness from such different directions we seem to share the same general outlook: the hard problem is a distraction; consciousness is not an added extra to everything else we do; and our false intuitions are the major stumbling block to escaping from dualism.

I can only thank Emily for making this third edition not only possible but, I think, the best yet.


EMILY

Sue had mentioned several times that she'd been asked to do a third edition but wasn't sure she could face it. I don't know quite why it was that on the third or fourth occasion, sometime in the summer of 2014, it occurred to me to offer to help. My academic background is in neither psychology nor neuroscience, nor even in philosophy, but in literary studies. But despite my predictable teenage rebellion against my psychologist parents, during my doctorate I'd found myself returning to the scientific fold by investigating the experience of reading Kafka,

- P R E F A C E S

and turning to lots of the same ideas Sue worked with – and even citing her quite often. And since then I’ve thought of myself as poised on the edges of many disciplines – quite a few of them the ones that make up this book.

I’d always thought this a wonderful, and surreally ambitious, book, and I hated the idea of it becoming gradually obsolete. Had I known quite how much time and energy the third edition would ask of me, or how hopeless the task would feel at times, I’m not completely sure I’d have made the offer. The process of co-authoring a book at all, let alone with my mother, let alone when living some of the time in her house, let alone when trying to do justice to the past six years of developments across all the fields that consciousness studies encompasses without adding many more words, has been something of an existential learning curve. Yet we’ve had lots of fun, too, and Sue has been very brave in letting me rip her baby to shreds and put it back together again – and now, three years later, it’s nearly over and I’m proud of what we’ve done: make an already great book, I think, even better.



Acknowledgements

FIRST EDITION

I would like to thank the following people who have helped me with arguments and discussion, who advised me on how to set about writing a textbook, or who have read and commented on parts of the manuscript. The very thorough reviewing process of Oxford University Press meant that I was able to improve the text in many ways as it went along. My thanks for help with the first edition to David Chalmers at Australian National University, John Crook, Dan Dennett at Tufts University, Stan Franklin at University of Memphis, David Goodworth, Nicky Hayes, Philip Merikle at University of Waterloo, Alva Noë at University of California, Berkeley, and Susan Schneider at the University of Pennsylvania.

SECOND EDITION

I would like to thank everyone who helped me with comments and suggestions for the second edition, including Paula Droege at Pennsylvania State University, Jay Gould at the University of West Florida, William Lycan at the University of North Carolina, Andrew Pessin at Connecticut College, Lisa Portmess at Gettysburg College, and Thomas Smythe at Carolina Central State. I would also like to thank the many referees as well as my colleagues, Guy Saunders and Jackie Andrade, and my agent Mandy Little.

THIRD EDITION

We are grateful to all those who helped shape this new edition, especially the anonymous readers who dedicated so much time and effort to reading the entire manuscript and commenting on it in detail. We were unable to act on all the excellent suggestions, but the final version is significantly stronger for this rich input. The inevitable mistakes and omissions that remain are our own. We thank Jackie Andrade at Plymouth University for helpful comments early on, Matt Tompkins for advice about magic, and Ilya Afanasyev, Felix Budelmann, and Chiara Cappellaro for help with translations. Our thanks go also to our editorial team at Routledge, including Liz Burton, Ceri McLardy, Holly Omand, and Sadé Lee; to Marie Roberts at Apex CoVantage; to our excellent copyeditor; and to Sue's agent Donald Winchester. Finally, we appreciate everything our partners have done – through patience, encouragement, cooking, and tea-making – to help keep us sane(ish).

The author and publishers would like to thank the following for permission to use the copyright material in this book.

Wellcome Library, London. Wellcome Images images@wellcome.ac.uk <http://wellcomeimages.org> Descartes: The Nervous System. Diagram of the brain and the pineal gland. De Homine Descartes, Rene Published: 1662 Copyrighted work available under Creative Commons Attribution only licence CC BY 4.0 for Figure 1.2; Jolyon Troscianko for Figure 10.8 and for the cartoons (see www.jolyon.co.uk); Nina Leen / Contributor / Getty Images for Figure 1.4; California Department of Fish and Wildlife / Flickr for Figure 2.1; 'Café Wall', Steven Battle, 2010, Shared under Creative Commons Attribution-Share Alike 3.0 Unported license for Figure 3.2; Wellcome Collection for Figure 3.3; Ramachandran and Blakeslee, 1998 for Figures 3.6, 3.7; Adam Hart-Davis for Figures 3.8, 7.5, 10.1, 10.9, 13.7; Mack, Arien, and Irvin Rock., Inattentional Blindness, Figure 4.13, p. 111, © 1998 Massachusetts Institute of Technology, by permission of The MIT Press for Figure 3.9; Metzinger, Thomas (ed.), Neural Correlates of Consciousness: Empirical and Conceptual Questions, Figure 15.2, p. 234, © 2000 Massachusetts Institute of Technology, by permission of The MIT Press for Figure 4.5; Ogawa et al. 2003 (Fig 3) for Figure 4.6; Redrawn from Schneider, 2009 for Figure 4.7; HarperCollins Publishers from VS Ramachandran and S Blakeslee, *Phantoms of the Brain* (Fourth estate, 1998) for Figure 4.9; Dennett, *Consciousness explained*, Little Brown 1991 for Figure 5.2; Shepard and Metzler, Rotation of three-dimensional objects, *Science*, 1971 for Figure 5.3; James Anderson for Figure 5.5; Baars, 1997a, p. 300 for Figure 5.6; Dehaene et al., Conscious, preconscious, and subliminal processing, *Trends in Cognitive Sciences*, 2006 for Figure 5.7; Penrose, 1994b for Figure 5.8; Redrawn from Crick, *The astonishing hypothesis*, Scribner's, 1994 for Figure 6.2; Redrawn from Engel et al., Temporal binding, binocular rivalry, and consciousness. *Consciousness and Cognition*, 1999 for Figure 6.3; Dover Publications for Figure 6.4; Luria, *The mind of a mnemonist: A little book about a vast memory*, Jonathan Cape, 1968 for Figure 6.6; Ramachandran and Hubbard, *Synaesthesia – A window into perception, thought and language*. *Journal of Consciousness Studies*, 2001 for Figure 6.7; Redrawn from Gazzaniga 1992 for Figure 6.9; Reprinted from Gazzaniga and LeDoux (1978), in Gazzaniga, 1992, p. 128 for Figure 6.10; Nature J Marshall and P Halligan, *Blindsight and insight in visuo-spatial neglect*, *Nature* 336, 766-77 (1988) for Figure 6.13 and 6.14; Evans et al., 2011 (Fig. 2) for Figure 7.3; Lutz et al., Attention regulation and monitoring in meditation. *Trends in Cognitive Sciences*, 2008 for Figure 7.8; adapted from Malinowski, 2013 for Figure 7.9; Popper and Eccles, *The self and its brain*, Springer, 1977 for Figure 8.2; René Descartes [Public domain], via Wikimedia Commons for Figure 8.3; Emmett Anderson / Wikimedia Melbourne Australian Open 2010 Venus Serve 5 for Figure 8.7; Castiello et al., Temporal dissociation of motor responses and subjective awareness: a study in normal subjects. *Brain*, 1991 for Figure 8.8; Milner and Goodale 1995 for Figure 8.9, 8.10, and 8.11; David Chalmers / Dover Publications for Figure 8.14; *Human volition: towards a neuroscience of will*, Patrick Haggard, 2008, Macmillan Publishers for Figure 9.1 and 9.2; Adapted from Libet et al., 1979 for Figure 9.3; *Journal of Consciousness Studies*, Libet, 1999, p. 51 for Figure 9.5; Eagleman and Holcombe, *Causality and the perception of time*, *Trends in Cognitive Science*, 2002 for Figure 9.6; Mary Evans Picture library for Figure 9.7; Redrawn from Wegner, *The illusion of conscious will*, MIT Press, 2002 for Figure 9.8; Miles 2011 for Figure 9.9; Adapted from Dennett, 1995 for Figure 10.2; Mary Evans Picture Library

for Figure 10.3; H. Zell, *Octopus vulgaris*, Octopodidae, Common Octopus; Staatliches Museum für Naturkunde Karlsruhe, Germany./ Wikimedia for Figure 10.5; *Journal of Consciousness Studies*, Sloman and Chrisley, 2003, p. 15 for Figure 10.6; Photographer / stevebloom.com for Figure 10.7; Osvaldo Cairo Battistutti, 2011 for Figure 10.10; Dr David Bygott for Figure 10.11; © Stuart Conway for Figure 10.13; Feinberg and Mallatt, *the ancient origins of consciousness*, MIT Press, 2016 for Figure 11.4; Redrawn from Humphrey, *The mind made flesh*, OUP, 2002 for Figure 11.5 and 11.8; Steven Mithen, 1996 for Figure 11.6; Barlow, *The biological role of consciousness*, in Blakemore and Greenfield, *Mindwaves*, Blackwell, 1987 for Figure 11.7; Sakurambo/Wikimedia Commons for Figure 11.9; Photo by DAVID ILIFF. License: CC-BY-SA 3.0 for Figure 11.11; akg-images for Figure 12.1; Biswarup Ganguly https://commons.wikimedia.org/wiki/File:Termite_Mound_-_Indian_Institute_of_Technology_Campus_-_Kharagpur_-_West_Midnapore_2013-01-26_3676.JPG for Figure 12.4; Topfoto.co.uk for Figure 12.5; Getty Images for Figure 12.6; newsoffice@mit.edu for Figure 12.9; Adapted from Aleksander 2005 for Figure 12.12; Owen Holland <owen@owenholland.org> for Figure 12.13; Redrawn from Holland, *A strongly embodied approach to machine consciousness*, *Journal of Consciousness Studies*, 2007 for Figure 12.14; Redrawn from Sloman and Chrisley, *Virtual machines and consciousness*, *Journal of Consciousness Studies*, 2003 for Figure 12.15; Luc Steels for Figure 12.16; Kędziński, J., Muszyński, R., Zoll, C. et al. *Int J of Soc Robotics* (2013) 5: 237 for Figure 12.17; PETER MENZEL/SCIENCE PHOTO LIBRARY for Figure 12.18; Wikipedia Commons for Figure 12.19; Adapted from Tart, *States of consciousness*, Dutton & Co, 1975 for Figure 13.3; Laureys, *The neural correlate of (un)awareness*, *Trends in Cognitive Sciences*, 2005 for Figure 13.4; Kirsch 2011, Fig. 1 (see Kirsch, *The altered state issue: Dead or alive?* *International Journal of Clinical and Experimental Hypnosis*, 2011) for Figure 13.9; Bettmann / Contributor / Getty Images for Figure 14.2; Werner Forman Archive/Private Collection for Figure 14.3; Siegel, *Hallucinations*, *Scientific American*, 1977 for Figure 14.4a; David Howard for Figure 14.4b and 15.5; Redrawn from Cowan, *Spontaneous symmetry breaking in large scale nervous activity*, *International Journal of Quantum Chemistry*, 1982 for Figure 14.5; Siegel and Jarvik, *Drug-induced hallucinations in animals and man*, in Siegel and West, *Hallucinations: Behavior, experience, and theory*, Wiley 1975 for Figure 14.6; Faith Goble Flickr for Figure 14.7; Hobson, 2002, Figure 2 for Figure 15.1; Redrawn from Hobson, *Dreaming*, OUP, 2002 for Figure 15.2; Redrawn from J. Allan Hobson in *Nature Reviews*, Vol. 10, November 2009 for Figure 15.3; Dr. Stephen LaBerge for Figures 15.7 and 15.8; Redrawn from Erlacher and Schredl, *Do REM (lucid) dreamed and executed actions share the same neural substrate?* *International Journal of Dream Research*, 2008 for Figure 15.9; Redrawn from Blanke and Arzy, *The out-of-body experience: Disturbed self-processing at the temporo-parietal junction*, *Neuroscientist*, 2005 for Figure 15.10; Carrington, *Modern psychical phenomena: Recent researches and speculations*, Kegan Paul, Trench, Trubner & Co., 1919 for Figure 15.13; R.A. Watters, from the Society for Psychical Research archives for Figure 15.16; Adapted from Blanke et al., *Stimulating illusory own-body perceptions*, *Nature*, 2002 for Figure 15.18; Adapted from Lenggenhager et al., *Video ergo sum: Manipulating bodily self-consciousness*, *Science*, 2007 and Ehrsson, H., *The experimental induction of out-of-body experiences*, *Science*, 2007. Muldoon and Carrington, 1929 for Figure 15.20; Adarsh Kumar / EyeEm / Getty Images for Figure 16.1; Paramount Pictures / Getty Images for Figure 16.2; Prince, *The*

dissociation of a personality Longmans, Green, and Co., 1906 for Figure 16.3; Sage Publications from R. Harre and G. Gillett, *The Discursive Mind* (Sage Publications) for Figure 16.4; Orion Publishing Group: cartoon by Daniel Postgate from *S. Law - The Philosophy Files* (Orion 2000) for Figure 16.5; Jason Hawkes / CORBIS for Figure 16.6; Jolyon Troscianko / EyeEm / Getty Images for Figure 16.10; FJ Varela *Neurophenomenology: A methodological remedy for the hard problem*. *Journal of Consciousness studies* 3(4), 330-49 for Figures 17.3 and 17.4; André Hatala [e.a.] (1997) *De eeuw van Rembrandt*, Bruxelles: Crédit communal de Belgique / Wiki Commons for Profile 1.1; Wiki Commons for Profile 1.2; David Chalmers for Profile 2.1; Patricia Churchland for Profile 2.2; V.S. Ramachandran for Profile 3.1; F. Imaoglu / Koch for Profile 4.1; Alonso Nichols, Tufts University / Daniel Dennett for Profile 5.1; Bernard Baars for Profile 5.2; Giulio Tononi for Profile 6.1; Robert Adam Mayer / Michael Graziano for Profile 7.1; Mel Goodale for Profile 8.1; Victor Albrow / Andy Clark for Profile 8.2; Wegner for Profile 9.1; Photo by Jana Lenzova / Richard Dawkins for Profile 10.1; Rosalie Winard / Temple Grandin for Profile 10.2; Nicholas Humphrey for Profile 11.1; Turing Archive for Profile 12.1; John Searle for Profile 12.2; Owen Holland for Profile 12.3; Courtesy of JGU Pressestelle for Profile 13.1; Ron Siegel for Profile 14.1; Allan Hobson for Profile 15.1; Antti Revonsuo for Profile 15.2; By Allan Ramsay - National Galleries Scotland, Public Domain, <https://commons.wikimedia.org/w/index.php?curid=51564591> for Profile 16.1; Joan Halifax/Upaya/Flickr for Profile 17.1; Sam Harris for Profile 18.1.

To help you get more out of this book visit:

www.routledge.com/cw/blackmore

Permissions granted:

Stephen Baxter, *Manifold: Time*. Copyright © Stephen Baxter 1999. First published in Great Britain by HarperVoyager. Published by permission of Stephen Baxter c/o Selectric Artists LLC.

Reprinted by permission from Macmillan Publishers Ltd: NATURE, T. Chiang, *What's Expected of Us*, 436(7047), 150, copyright 2005.

THE POEMS OF EMILY DICKINSON: READING EDITION, edited by Ralph W. Franklin, Cambridge, Mass.: The Belknap Press of Harvard University Press, Copyright © 1998, 1999 by the President and Fellows of Harvard College. Copyright © 1951, 1955 by the President and Fellows of Harvard College. Copyright © renewed 1979, 1983 by the President and Fellows of Harvard College. Copyright © 1914, 1919, 1924, 1929, 1930, 1932, 1935, 1937, 1942 by Martha Dickinson Bianchi. Copyright © 1952, 1957, 1958, 1963, 1965 by Mary L. Hampson.

From *The Magus* by John Fowles published by Jonathan Cape. Reproduced by permission of The Random House Group Ltd. ©1966.

From *The French Lieutenant's Woman* by John Fowles published by Jonathan Cape. Reproduced by permission of The Random House Group Ltd. ©1969.


Thomas Mann, *Der Tod in Venedig*. In: ders., *Gesammelte Werke in dreizehn Bänden*. Band VIII. Erzählungen. © S.Fischer Verlag GmbH, Frankfurt am Main, 1960, 1974.

Saturday by Ian McEwan. Published by Jonathan Cape, 2005. Copyright © Ian McEwan. Reproduced by permission of the author c/o Rogers, Coleridge & White Ltd., 20 Powis Mews, London W11 1JN.

We thank Peter Watts for his kind permission to use three passages from *Blind-sight*, published by Tom Doherty, Tor, New York, 2006.

We thank Robert Eno for his kind permission to use a passage from his translation of *Zhuangzi: The Inner Chapters* (2010/2016), which you can find online at <http://www.indiana.edu/~p374/Zhuangzi.pdf>.

We thank the Writers House agency for their kind permission to use an excerpt from Hermann Hesse's *Das Glasperlenspiel* (The Glass Bead Game), 1943.



Introduction

WELCOME PERPLEXITY

If you think you have a solution to the problem of consciousness, you haven't understood the problem. That's not strictly true, of course. You may either be a genius and have found a real solution, or be sufficiently clear-sighted to understand why there was no problem in the first place. More likely, however, is that you are falling into a number of tempting traps that help you evade the real issues.

The American philosopher Thomas Nagel once observed that 'Certain forms of perplexity – for example, about freedom, knowledge, and the meaning of life – seem to me to embody more insight than any of the supposed solutions to those problems' (1986, p. 4). This may be equally true of the problem of consciousness. Indeed, the puzzlement can be part of the pleasure, as philosopher Colin McGinn points out: 'the more we struggle the more tightly we feel trapped in perplexity. I am grateful for all that thrashing and wriggling' (1999, p. xiii).

If you want to think about consciousness, confusion is necessary: mind-boggling, brain-hurting, *I can't bear to think about this stupid problem any more* confusion. For this reason, a great deal of this book is aimed at increasing your perplexity rather than reducing it. So, if you do not wish your brain to hurt (though of course strictly speaking brains cannot hurt because they do not have any pain receptors – and, come to think of it, if your toe, which does have pain receptors, hurts, is it really your toe that is hurting?), stop reading now and choose a more tractable problem to study.

Our motivation for wishing to stir up perplexity is not cruelty or cussedness, nor the misplaced conviction that long words and difficult arguments are signs of cleverness or academic worth. Indeed, we think the reverse: that the more difficult a problem is, the more important it becomes to use the simplest words and sentences possible. So, we will try to keep our arguments as clear and simple as we can while tackling what is, intrinsically, a very tricky problem.

Part of the problem is that 'consciousness' has no generally accepted definition in either science or philosophy despite many attempts to define it (Nunn, 2009). The word is common enough in everyday language, but is used in different ways. For example, 'conscious' is often contrasted with 'unconscious', and



PROFILE 0.1

Susan Blackmore (b. 1951)



As a student in Oxford, reading physiology and psychology, Sue Blackmore had a dramatic out-of-body experience which convinced her that consciousness could leave the body and made her determined, against much sound advice, to study parapsychology. She learned to read Tarot cards, sat with mediums, and trained as a witch, but her 1979 PhD thesis contained only numerous failed experiments on extra-sensory perception and psychokinesis. Becoming ever more sceptical of paranormal claims, she turned to studying the experiences that foster paranormal belief, including near-death experiences, sleep paralysis, and dreams, eventually concluding that parapsychology is a red herring in any attempt to understand consciousness. Meditation proved far more helpful, and she has been practising Zen since the early 1980s. She carried out one of the first experiments on change blindness, and her books include the controversial bestseller *The Meme Machine* as well as books on OBEs, NDEs, meditation, and consciousness. While at the University of the West of England in Bristol, she taught the consciousness course on which this book is based, but finally decided that the only way to learn more about consciousness was to give up the job and write this book. Since then she has been a freelance writer and lecturer and is now working (again) on out-of-body experiences, tremes (technological memes), and (unsuccessful) children's books. She plays in a samba band, loves painting, kayaking, and her garden, and is learning powerlifting. She is Visiting Professor in Psychology at the University of Plymouth.

is taken as more or less equivalent to 'responsive' or 'awake'. 'Conscious' is also used to mean the equivalent of knowing something, or attending to or being aware of something, as in 'She wasn't conscious of the embarrassment she'd caused' or 'He wasn't conscious of the rat creeping up quietly under his desk'. Different theories emphasise different aspects of what we might mean by consciousness, but the term is most broadly used to mean the equivalent of 'subjectivity' or personal experience, and this is the sense in which it is used throughout this book.

Another problem is that consciousness studies is a relatively new and profoundly multidisciplinary subject. This means we can draw on a rich variety of ideas from neuroscience, philosophy, psychology, biology, and other fields, but it can also make life difficult because people from these different disciplines sometimes use the same words in completely different ways. Students of psychology are our primary audience in this book, but we have tried to cover all of the major approaches in consciousness studies, including psychology, philosophy, artificial intelligence, neuroscience, and first- and second-person methods, as well as 'non-traditional' approaches centred on spirituality or 'altered states' of consciousness. We have also included excerpts from novels, stories, poems, diaries, and letters to help you explore consciousness with the help of a wider range of great writers and thinkers. Our emphasis is on a science of consciousness based on empirical findings and testable theories, but there are many forms this science can take. Throughout the book we will be confronted by questions about how the nature of consciousness (its ontology) is related to the possibility of gaining knowledge about it (the epistemology) and the methods we choose to do so (the methodology). We have no easy answers, other than to keep reminding you (and ourselves) that there is no such thing as a neutral question or method. Even the ordinary language we use to think with pushes us in one direction or another from the very outset.

No single existing method of studying consciousness has all the answers. Because the brain is the most complicated organ in the human body, it is easy to think that it must hold the answer to the mystery of consciousness. But when people have tried to fit consciousness neatly into the usual ways of doing brain science, they find they cannot do it. This suggests that somewhere along the line we are making a fundamental mistake or relying on some false assumptions. Rooting out one's prior assumptions is never easy and can be painful. But that is probably what we have to do if we are to think clearly about consciousness.

THE ORGANISATION OF THE BOOK

This book is divided into six relatively independent sections containing three chapters each. Each section is designed to stand alone, for use as the topic for a lecture, or several lectures, or to be read independently as an overview of the area. However, all of them depend on the ideas outlined in Section One, so if you choose to read only parts of the book, we would recommend starting with Section One, on the nature of the problem.

There is an accompanying website at www.routledge.com/cw/blackmore. This provides a complete list of references with weblinks where possible, suggested questions for class or self-assessment, and further information, demos, and audio-visual materials, as well as updates to the printed book. It also provides some suggestions of different ways you can navigate the book depending on your specific interests.

Each chapter contains not only a core text, but also profiles of selected authors, explanations of key concepts, exercises to do on your own, and suggestions for activities and discussions that can be done in groups.

At the end of each chapter is a list of suggested readings with brief descriptions. The readings are chosen to be short and readily accessible and to give a quick way into each topic. They should also be suitable as set reading between lectures for those whose courses are built around the book. For each chapter we include at least one reading (highlighted in red) which offers multiple perspectives on a topic, whether through peer commentaries on a target article, a range of views on a question or concept, or case studies; these may be useful as the basis for seminar discussions.

Each chapter includes one or more quotations from literary works highlighted in orange. Many of them come from famous writers, and you may know some of them already. We hope they will do two things: on the one hand, enrich your understanding of the often strange ideas about consciousness that we will be encountering; and on the other, enhance your appreciation of the authors and works we quote from by revealing the links between the ideas they have long been exploring and the problems that contemporary psychology, philosophy, and neuroscience are still battling with. Many originate in languages other than English, and we have provided the most faithful translations we could. This may also help you think about how different languages offer tools for thinking about consciousness.

We also provide shorter quotes in the margins, often repeated from the main text. Our advice is to learn those that appeal to you by heart. Rote learning seems hard if you are not in the habit, but it gets quickly easier with practice. Having quotations at your mental fingertips looks most impressive in essays and exams but, much more important, it provides a wonderful tool for thinking with. If you are walking along the road or lying in bed at night, wondering whether there really is a 'hard problem' or not, your thinking will go much better if you can bring instantly to mind Chalmers's definition of the problem, or the exact words of his major critics. Often a short sentence is all you need to get to the crux of an argument and criticise it: what assumptions underlie it, and what exactly does it help you to understand better?



PROFILE 0.2

Emily Troscianko (b. 1982)



Emily is Sue's daughter and has many (mostly fond) childhood memories of Sue's strange explorations of the paranormal, alien abductions, and memes, as well as of morning meditation sessions together before school. Emily studied French and German as an undergraduate at Oxford, and stayed there to do a doctorate on the works of Franz Kafka. Asking the question 'Why is Kafka's writing so powerful?' led her to investigate theories of vision, imagination, and emotion, and to conduct her own experiments on how readers respond to different kinds of fictional text. Having suffered from anorexia from age 16 to 26, she later began to connect her interest in mental health with her understanding of literary reading, starting to explore how fiction-reading might have effects on mental illness, and vice versa. Her current work is a mixture of cognitive-literary and medical-humanities research and various kinds of writing for audiences beyond academia. Like Sue, she seems to have had to give up having a job to write this book. When not writing, she can often be found driving her cow-spotted campervan around Britain, captaining her narrowboat along the Thames, or lifting heavy things (sometimes with Sue) in a powerlifting gym.

PUTTING IN THE PRACTICE

Consciousness is a topic like no other. Right now, this very minute, you are probably convinced that you are conscious – that you have your own private experience of the world – that you are personally aware of things going on around you and of your own inner states and thoughts – that you are inhabiting your own private world of awareness – that there is something it is like to be you. This is what is meant by being conscious. Consciousness is our first-person view on the world.

In most of our science and other academic studies, we are concerned with third-person views – with things that can be verified by others and agreed upon (or not) by everyone. But what makes consciousness so interesting is that it cannot be agreed upon in this way. It seems private. It seems like something on the inside. I cannot know what it is like to be you. And you cannot know what it is like to be me.

So, what is it like to be you? What are you conscious of now?

Well...? Take a look. Go on. Really. Take a look and try to answer the question 'What am I conscious of now?'

Is there an answer? If there is an answer, you should be able to look and see. You should be able to tell someone else, or at least know for yourself, what you are conscious of now, and now, and now – what is 'in' your stream of consciousness. If there is no answer, then our confusion must be very deep indeed, for it certainly seems as though there must be an answer – that I really am conscious right now, and that I am conscious of some things and not others. If there is no answer, then at the very least we ought to be able to understand why it feels as though there is.

So, take a look and first decide whether there is an answer or not. Can you do this? You will probably decide that there is: that you really are conscious now, and that you are conscious of some things and not others – only it is a bit tricky to see exactly what

this is like, because it keeps on changing. Every time you look things have moved on. The sound of the hammering outside that you were conscious of a moment ago is still going on but has changed. A bird has just flitted past the window, casting a brief shadow across the window sill. Oh, but does that count? By the time you asked the question 'What am I conscious of now?'; the bird and its shadow had gone and were only memories. But you were conscious of the memories, weren't you? So maybe this does count as 'what I am conscious of now' (or, rather, what I was conscious of then).

The morning was hot, and the exercise of reading left her mind contracting and expanding like the main-spring of a clock, and the small noises of midday, which one can ascribe to no definite cause, in a regular rhythm. It was all very real, very big, very impersonal, and after a moment or two she began to raise her first finger and to let it fall on

the arm of her chair so as to bring back to herself some consciousness of her own existence. She was next overcome by the unspeakable queerness of the fact that she should be sitting in an arm-chair, in the morning, in the middle of the world. Who were the people moving in the house – moving things from one place to another? And life, what was that? It was only a light passing over the surface and vanishing, as in time she would vanish, though the furniture in the room would remain. Her dissolution became so complete that she could not raise her finger any more, and sat perfectly still, listening and looking always at the same spot. It became stranger and stranger. She was overcome with awe that things should exist at all . . . She forgot that she had any fingers to raise . . . The things that existed were so immense and so desolate . . . She continued to be conscious of these vast masses of substance for a long stretch of time, the clock still ticking in the midst of the universal silence.

(Virginia Woolf, *The Voyage Out*, 1915)

You will probably find that if you try to answer the first question, many more will pop up. You may find yourself asking ‘How long is “now”?’, ‘Was I conscious before I asked the question?’, ‘Who is asking the question?’, ‘What does it mean to “look” “inside”?’ Indeed, you may have been asking such questions for much of your life. Teenagers commonly ask themselves difficult questions like these and don’t find easy answers. Some go on to become scientists or philosophers or meditators and pursue the questions in their own ways. Many just give up because they receive no encouragement, or because the task is too difficult. Nevertheless, these are precisely the kinds of questions that matter for studying consciousness. That is why each chapter includes a ‘practice’ task with a question to work on in between your reading.

Every question and every practice takes only one angle on the problem of consciousness. Some – including the one we started with here – may not be helpful for you. But we hope that cumulatively, day by day, they will help you. One of us, Sue, has been asking questions like these many times a day for about thirty years, often for hours at a stretch. She has also taught courses on the psychology of consciousness for more than ten years and encouraged her students to practise asking these questions. Over the years she has learned which ones work best, which are too difficult, in which order they can most easily be tackled, and how to help students who get into a muddle with them. And Emily has come to puzzle over consciousness from different starting points – from questions about how we experience fictional worlds to questions about what it means to be mentally healthy or ill. We encourage you to work hard, not just at the science but at your own personal practice, alone and together with others who are questioning, too.

GETTING THE BALANCE RIGHT

A lot of this book is about so-called third-person views. You will learn about neuroscientific experiments, philosophical inquiries, and psychological theories. You will learn to be critical of theories of consciousness, and of the many ways

of testing one against another. But underlying all of this is the first-person view, which is what it's all about. Some scientists and philosophers try to connect the two; some create bridges between the first and the third person by thinking about the 'second person', or how 'my' experience is already shaped by other people. Still, the distinction between more theoretical and more personal ways of studying consciousness remains, and you must strike a balance between them.

That balance will be different for each of you. Some will enjoy the self-examination and find the science and philosophy hard. Others will lap up the science and find the personal inquiry troubling or trivial. However it is for you, remember that both are needed, and you must find your own balance between them. To those who object that self-questioning is a waste of time or even 'childish', we can only say this: since we are studying subjective experience, we must have the courage to become familiar with subjective experience.

As you become acquainted with the growing literature of consciousness studies, and if you have managed to strike a balance between the work of observing your own experience and the work of explaining it, you will begin to recognise those writers who have not. At one extreme are theorists who say they are talking about consciousness when they are not. They may sound terribly clever, but you will soon recognise that they have never attended to their own experience. What they say simply misses the point. At the other extreme are those who waffle on about the meaning of inner worlds or the ineffable power of consciousness while falling into the most obvious of logical traps – traps that you will instantly identify and be able to avoid. Once you can spot these two types, you will save a lot of time by not struggling with their writings. There is so much to read on the topic of consciousness that finding the right things to struggle with is quite an art. We hope this book will help you to find reading that is worthwhile for you, and to avoid the time-wasting junk. We cannot claim to have been completely impartial, but we have tried to be your sceptical guides through this difficult field, to help you find your own way through it.

WARNING

Studying consciousness will change your life. At least, if you study it deeply and thoroughly it will. As the American philosopher Daniel Dennett says, 'When we understand consciousness – when there is no more mystery – consciousness will be different' (1991, p. 25). None of us can expect to thoroughly 'understand consciousness'. It is still not even clear what that would mean. Nonetheless, we do know that when people really struggle with the topic, they find that their own experience, and their sense of self, change in the process.

These changes can be uncomfortable. For example, you may find that once-solid boundaries between the real and unreal, or the self and other, or humans and other animals or robots, or you right now and someone in a coma, begin to look less solid. You may find that your own certainties – about the world out there, or ways of knowing about it – seem less certain. You may even find yourself beginning to doubt your own existence. Perhaps it helps to know that many people have had these doubts and confusions before you, and have survived.

The difficulties I have in talking to people, which others must find incredible, come from the fact that my thinking, or rather the content

'The most beautiful thing we can experience is the mysterious. It is the source of all true art and science.'

(Einstein, 1930)

of my consciousness, is quite foggy, that as far as it concerns only myself I rest in it untroubled, sometimes even self-satisfied, but that human conversation requires pointedness, stability, and sustained coherence, things that do not exist in me. No one will want to lie in clouds of fog with me, and even if someone did, I cannot drive the fog out of my head; between two people it melts away and is nothing.”


(Franz Kafka (1990), diary entry, 24 January 1915, our translation)

Indeed, many would say that life is easier and happier once you get rid of some of the false assumptions we so easily tend to pick up along the way. But that is for you to decide for yourself. If you get into difficulties, we hope you will be able to find appropriate help and support from peers, teachers, or other professionals. If you are teaching a course using this book, you should be prepared to offer – and seek out – that support yourself, or be able to advise students on how to find help when they need it.

Some of Sue’s classes included a few students who held religious convictions or believed in God. They usually found that these beliefs were seriously challenged by the course. Some found this difficult, for example because of the role of faith in family ties and friendships, or because their beliefs gave them comfort in the face of suffering and death, or because religion provided a framework for thinking about self, consciousness, and morality in terms of a spirit or soul. So, if you do have such beliefs you should expect to find yourself questioning them. It is not possible to study the nature of self and consciousness while labelling God, the soul, the spirit, or life after death ‘off limits’.

Every year she taught courses on consciousness, Sue gave this same warning to students – both in person and in writing. Every year, sooner or later, one of them came to her, saying ‘You never told me that. . .’. Happily, most of the changes are, in the end, positive, and the students are glad to have been through them. Even so, we can only repeat our warning and hope that you will take it seriously. **Studying consciousness will change your life.** Have fun.

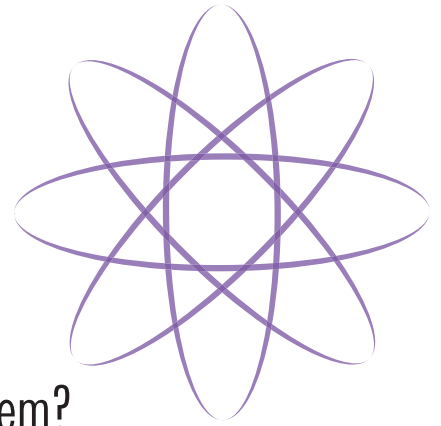
‘Warning – studying consciousness will change your life’.



The problem

SECTION

ONE



What's the problem?

C
H
A
P
T
E
R

ONE

WHAT IS THE WORLD MADE OF?

The problem of consciousness is related to some of the oldest questions in philosophy: What is the world made of? How did it get here? Who or what am I? What is the point of it all? In particular, it is related to the mind–body problem – that is, what is the relationship between the physical and the mental?

In the early twenty-first century, many people use the term ‘consciousness’ quite unproblematically in everyday language to refer to their own private experience or awareness. It is no longer synonymous with ‘mind’, which has many other meanings and uses, and which seems to have lost some of its mystery. This is mainly because we are rapidly learning how the brain works. We know about the effects of brain damage and drugs, about neurotransmitters and neuromodulators, about how changes in the firing of brain cells accompany changes in a person’s experience, and about how all this relates to the rest of the nervous system and the body. We might expect all this knowledge to have clarified the nature and causes of conscious awareness, but it doesn’t seem to have done so. Consciousness remains a mystery.

In many other areas of science, increasing knowledge has made old philosophical questions obsolete. For example, no one now agonises over the question ‘what is life?’ The old theories of a ‘vital spirit’ or *élan vital* are superfluous when you understand how biological processes make living things out of non-living matter. As Daniel Dennett puts it, ‘the recursive intricacies of the reproductive machinery of DNA make *élan vital* about as interesting as Superman’s dread kryptonite’ (1991, p. 25). The point is not that we now know what *élan vital* is, but that we don’t

‘There is nothing that we know more intimately than conscious experience, but there is nothing that is harder to explain’.

(Chalmers, 1995a, p. 200)




FIGURE 1.1

'There exists no accepted definition of consciousness.'

(Dietrich, 2007, p. 5)

need it any more, because we know there is no such thing. The same is true of the 'caloric fluid' which was once needed to explain the nature of heat. Now that we think of heat as a form of energy, and know how various types of energy are transformed into each other, we know that the term 'caloric fluid' does not refer to anything which really exists.

Might the same happen with consciousness? The American philosopher Patricia Churchland thinks so, arguing that when our framework for understanding consciousness has evolved, consciousness 'may have gone the way of "caloric fluid" or "vital spirit"' (1988, p. 301).

Maybe it will. But so far it has not. And Chalmers says that we would be foolish to expect it to, since the *élan vital* was proposed as a way of explaining something else (how life is created from matter), and so could be discarded when we found a better explanation, whereas consciousness is something that itself needs to be explained. 'Experience is not an explanatory posit but an explanandum in its own right, and so is not a candidate for this sort of elimination' (1995a, p. 209). So maybe we should not expect this kind of cut-and-replace fix when it comes to consciousness. Indeed, the more we learn about the brain and behaviour, the more obviously difficult the problem of consciousness seems to be.

In essence, it is this. Whichever way we try to wriggle out of it, in our everyday language or in our scientific and philosophical thinking, we seem to end up with some kind of impossible dualism. Whether it is spirit and matter, or mind and brain; whether it is inner and outer, or subjective and objective – we seem to end up talking about two incompatible kinds of stuff. Maybe we are all 'natural-born dualists' (Bloom, 2004, p. xiii) with inescapable 'intuitions of distinctness' about mind and matter (Papineau, 2002). But you may disagree. You may, for example, say that you are a materialist – that you think there is only one kind of stuff in the world and that mind is simply the workings of that stuff. Problem solved. Yet if you take this line, or adopt many other popular ways of tackling the problem, you will only find that in thinking about consciousness, the dualism pops up somewhere else. Let's take an example.

Pick some simple object you have to hand and take a good look at it. You might choose a chair or table, the cat curled up on your desk, or a book. Anything will do. Let's take a pencil. You can pick it up, turn it round, play with it, write with it, put it down in front of you. Now ask yourself some basic questions. What do you think it is made of? What will happen if you hold it two feet above the floor and let go? If you leave the room and come back will it still be here?

Now think about your experience of the pencil. You may have felt its sharp point and texture, smelled its distinctive smell when you sharpened it, seen its colour and shape, and written with it. These experiences are yours alone. When you hold

the pencil at arm's length, you see the pencil from your own unique perspective. No one else can have exactly the same pencil-watching experience as you are having now. And what about the colour? How do you know that the way you see that yellow paint would be the same for someone else? You don't. This is what we mean by consciousness. It is your experience. No one else can know what it is like. No one else can get it from you. You can try to tell them, but words can never quite capture what it is like for you to be holding that pencil right now.

So where has this got us? It has forced us into thinking about the world in two completely different ways. On the one hand, there is your apparently private and intimately known experience of holding the pencil, and on the other, there is the real pencil out there in the world. How can your sensations be related to real existing objects in space? Does the activity in the visual cortex of your brain *cause* your experience of pencil-watching? If so, how? What makes the smell like *this* for you?

Probably everyone has a different sticking point on this. For Sue, it is this. I find that I have to believe both in subjective experiences (because I seem unquestionably to have them) and an objective world (because otherwise I cannot possibly explain why the pencil will drop when I let go, will still be here when I get back, or why you and I can agree that it is blunt and needs sharpening). For Emily, it is this. What could it possibly be that makes there be an experience of holding a pencil at all, rather than all the skin cells and nerve endings and muscular contractions being just like the pencil seems to be – 'dark inside'? Even with all our understanding of how the brain and the rest of the body works, we cannot understand how the subjective, ineffable thisness of experience arises from an objective world of actual pencils and living brain cells. These subjective and objective worlds seem to be too different from each other to be related at all. These are our own versions of the problem of consciousness – our own sticking points. You should look hard at the pencil and find out where yours lies.



PRACTICE 1.1

AM I CONSCIOUS NOW?

For this first exercise, we will give you more detailed guidance than for future ones. All the rest build on the same foundation, so you should find that if you practise this one frequently all the others will be easier.

The task is simply this.

As many times as you can, every day, ask yourself '**Am I conscious now?**'

The idea is not simply to provide an answer – for example, 'Yes' – twenty or a hundred times a day, but to begin looking into your own mind. When do you answer 'Yes'? and when 'No'? What does your answer mean?

You might like to ask the question and then just hold it for a little while, observing being conscious now. Since this whole book is about consciousness, this exercise is simply intended to get you to *look at, feel,*

listen to, smell, and taste what consciousness is, as well as to think and argue about it intellectually.

This sounds easy, but it is not. Try it and see. After a day of practising, or – if you are working through the book, before you go on to the next chapter – make notes on the following.

How many times did you do the practice?

What happened?

Did you find yourself asking other questions as well? If so, what were they?

Was it difficult to remember to do it? If so, why do you think this is?

You may have found that you had intended to do the practice but then forgot. If you need reminding you might try these simple tricks:

Ask the question whenever you hear or read the word ‘consciousness’.

Always ask the question when you go to the toilet.

Write the question on stickers and place them around your home or office.

Set a reminder on your phone.

Pair up with a friend to help you remind each other.

These cues may help, but you may still forget, which is odd because there is no very good excuse. After all, this little practice does not take up valuable time when you could be writing another essay, reading another paper, or struggling with a difficult argument. You can ask the question while doing any of these things; while walking along or waiting for the bus; while washing up or cooking; while cleaning your teeth or listening to music. You just keep on doing it, pose the question, and watch for a moment or two.

You must be interested in consciousness to be reading this book. So why is it so hard just to look at your own consciousness?

Are you conscious now?

They might not seem to, but these sticking points matter to some of the most difficult questions we face, like: whether to withdraw life support from someone who cannot move or respond to communication; how to treat non-human animals, unborn foetuses, and artificial intelligences; whether it makes sense to hold ourselves morally accountable for our actions; what drug-induced experiences and mental illness have in common (if anything); and how to do science. Your sticking point is fundamental to all the questions you have or haven’t ever asked about what makes you you.

CONSCIOUSNESS IN PHILOSOPHY

Philosophers over the millennia have struggled with versions of the problem of consciousness. Their solutions can be roughly divided into monist theories – which assert that there is only one kind of stuff in the world – and dualist theories, which propose two kinds of stuff.

For most people, dualism is the starting point. Many of our most natural ways of talking about ourselves, from ‘I need to get a grip on myself’ to ‘I nearly jumped out of my skin’, make dualism the default position. Ordinary language makes it hard to avoid separating a mysterious *myself* off from ‘my body’ and even from ‘my mind’: after all, if they are *mine*, then I cannot *be* them.

The best-known version of dualism is that of the seventeenth-century French philosopher René Descartes, and is therefore called Cartesian dualism. Descartes wanted to base his philosophy only on firm foundations that were beyond doubt. If he had been holding your pencil, he might have made himself imagine that it did not exist and that his senses were deceiving him, or that he was only dreaming, or even that an evil demon was systematically trying to fool him. But, he argued, in a famous passage in *Discourse on Method* (1637/1649), even if we doubt everything, there is still something that remains. The fact that he, Descartes, was thinking about this was proof that he, the thinker, existed. In this way he came to his famous dictum ‘je pense, donc je suis’, ‘I think, therefore I am’, and he called it ‘the first principle of the Philosophy I sought’ (pp. 51–52). In his later *Meditations on First Philosophy* (or just *Meditations*, 1641), Descartes concluded that this thinking self was not material, like the physical body that moves about mechanically and takes up space. In his view, the world consists of two different kinds of stuff: the extended stuff of which physical bodies are made, and the unextended, thinking stuff of which minds are made.

Descartes’s theory is a form of substance dualism. It can be contrasted with property dualism or dual aspect theory (which are sometimes considered to be forms of ‘anomalous monism’). According to property dualism, the world is composed only of one kind of substance (the physical kind), but can be described using mental terms or physical terms, even though one description cannot be reduced to the other. So, for example, if you are in pain, this fact can be described in mental terms, such as how it feels to you, or in physical terms, such as which sorts of neurons are firing where in your nervous system. This theory avoids the need for two different substances, but leaves open many questions about the relationship between the physical and mental properties, and therefore comes in many different versions.

The insuperable problem for substance dualism is how the mind interacts with the body when the two are made of different substances. For the whole theory to work, the interaction has to be in both directions. Physical events in the world and the brain must somehow give rise to experiences of that world – to thoughts, images, decisions, longings, and all the other contents of our mental life. In the other direction, thoughts and feelings must be able to influence the physical stuff. How could either of these work? Descartes supposed that the two interact through the pineal gland in the centre of the brain, but proposing a place where it happens does not solve the mystery. If thoughts can affect brain cells, then either



PROFILE 1.1

René Descartes (1596–1650)



Descartes was born near Tours in France, was educated at a Jesuit college, and was a staunch believer in an omnipotent and benevolent God. On 11 November 1619, he had a series of dreams which inspired him with the idea of a completely new philosophical and scientific system based on mechanical principles. He was not only a great philosopher, now often called ‘the father of modern philosophy’, but also a physicist, a physiologist, and a mathematician. He was the first to draw graphs and invented Cartesian coordinates, which remain a central concept in mathematics. He is best known for his saying ‘I think, therefore I am’ (*je pense, donc je suis*, or *cogito ergo sum*), which he arrived at using his ‘method of doubt’. He tried to reject everything that could be doubted and accept only that which was beyond doubt, which brought him to the fact that he, himself, was doubting. He described the human body entirely as a machine made of ‘extended substance’ (in the Latin, *res extensa*), but concluded that the mind, spirit, or soul (which he called the *animus*) must be a separate entity made of a non-spatial and indivisible ‘thinking substance’ (*res cogitans*). The two substances were connected through the pineal gland. This theory became known as Cartesian dualism, a term which is now used synonymously with substance dualism — that is, any theory that posits causal interactions between fundamentally distinct substances, material and immaterial. For the last twenty years of his life, Descartes lived mostly in Holland. He died of pneumonia in Sweden in 1650.

they work by magic or they must be using some kind of energy or matter. In this case, they are also physical stuff and not purely mental.

Do you see this egg? With this you can overthrow all the schools of theology, all the temples of the earth. What is this egg? [. . .] First there’s a dot that quivers, a thread that stretches and takes on colour, tissue that is formed; a beak, tips of wings, eyes, feet that appear; a yellowish matter that uncoils and produces intestines; it is an animal. This animal moves, writhes, cries out; I hear its cries through the shell; it clads itself with down; it sees [. . .] it has all your ailments; all your actions, it performs them. Will you claim, with Descartes, that it is a pure imitative machine? But little children will laugh at you, and philosophers will reply that if this is a machine then you are too. If you admit that between the animal and yourself there is no difference but in organisation, you will be showing good sense and reason, you will be honest; but from this people will conclude against you that from inert matter, arranged in a certain way, impregnated with other inert matter, with heat and motion, there results the capacity for sensation, life, memory, consciousness, passion, thought.

(Denis Diderot, *Conversation between d’Alembert and Diderot* [Entretien entre d’Alembert et Diderot], 1769, our translation)

Substance dualism does not work. Almost all contemporary scientists and philosophers agree on this. In 1949 the British philosopher Gilbert Ryle derided dualism as ‘the dogma of the Ghost in the Machine’ (p. 26), a phrase that has entered into common parlance. He said we should avoid phrases like ‘in the head’ and ‘in the mind’ because they make us think of minds as ‘queer “places”, the occupants of which are special-status phantasms’ (1949, p. 40). Ryle was influenced by Wittgenstein’s ‘ordinary-language philosophy’, which proposed that many philosophical problems are caused by misuses of language. Ryle argued that because we don’t know how to talk about mind, we often talk about it using the language of material cause and effect, but in the negative: we say, ‘Minds are not bits of clockwork, they are just bits of not-clockwork’ (p. 9). If we do this, we are making a category mistake: we are putting mind into the wrong category of




FIGURE 1.2 • According to Descartes, the physical brain worked by the flow of animal spirits through its cavities. The immaterial soul was connected to the body and brain through the pineal gland (H) which lies in the midline.

things, giving something non-material properties which logically and grammatically apply only to material things. Ryle did not reduce mental processes to physical processes; he tried to find a middle way between dualism and behaviourism – between the two mistakes of claiming, for instance, that saying is doing one thing and thinking is doing another, or that saying and thinking are the same thing. For Ryle, behaviours are not caused by mysterious mental states, and many mental states are best understood simply as dispositions to behave.

The view of mind as about *doing* rather than *being* is apparent in many modern descriptions of mind and self: ‘Minds are simply what brains do’ (Minsky, 1986, p. 287); ‘“Mind” is designer language for the functions that the brain carries out’ (Claxton, 1994, p. 37); and self is ‘not what the brain *is*, but what it *does*’ (Feinberg, 2009, p. xxi). Such descriptions make it possible to talk about some mental activities and mental abilities without supposing that there is a separate mind. This is probably how most psychologists and neuroscientists think of ‘mind’ today, and they do not agonise about what ‘mind’ really is.

Some psychologists and philosophers, especially those interested in the limitations of a brain-centric view, carefully avoid reducing the activity of the mind to solely neural activity. And some, like the American philosopher Alva Noë, do the same for consciousness itself: ‘Consciousness is not something that happens in us. It is something we do’ (2009, p. 160). With consciousness there is generally much less agreement, probably because a lot of the questions people used to ask about mind and self (or even the soul) are now being directed at ‘consciousness’ instead. But as we will see, how we think about any of the three will have important consequences for how we think about the other two. And central to them all is the thorny question of what the relationship is between mind and matter.




FIGURE 1.3 • Gilbert Ryle (1949) dubbed the Cartesian view of mind ‘the dogma of the Ghost in the Machine’.

‘[My position] is an innocent version of dualism, entirely compatible with the scientific view of the world’

(Chalmers, 2007, p. 360)

‘accepting dualism is giving up’

(Dennett, 1991, p. 37)

The twentieth century saw two notable attempts to make dualism work. In the 1970s, the philosopher of science Karl Popper and neurophysiologist John Eccles (1977) proposed a theory of dualist interactionism. They argued that the critical processes in the synapses of the brain are so finely poised that they can be influenced by a non-physical thinking and feeling self. Thus the self really does control its brain (Eccles, 1994). How it does so, they admit, remains mysterious. The physiologist Benjamin Libet (2004) proposed that a non-physical ‘conscious mental field’ is responsible for the unity and continuity of subjective experience and for free will. Somewhat like a physical force field, it emerges from brain activity, but it can then communicate within the cerebral cortex without using the neural connections and pathways. But how it does this he does not explain.

More recently, the Australian philosopher David Chalmers has proposed a ‘naturalistic dualism’, which he calls ‘an innocent form of dualism, entirely compatible with the scientific view of the world’ (2007, p. 360). Rather than contradicting physical principles, it suggests that new ‘bridging principles’, in the form of psychophysical laws, are needed to explain how experience arises from physical processes even though the physical world is causally closed. The theory is a version of property dualism or dual aspect theory, with the central concept of *information* taking both phenomenal and physical forms. As in other versions of dualism, however, the bridge arguably does not reach the whole way across the gap.

Because hardly anyone admits to being a dualist any more, but dualism is so hard to get away from, the philosopher Daniel Dennett (1991) coined the term ‘Cartesian materialism’ to describe the position of pretending to be a materialist but relying on dualist concepts – particularly the idea that there is an identifiable time and place where everything comes together in the brain and ‘consciousness happens’. Dennett’s PhD was supervised by Ryle, and he shares Ryle’s view of the importance of paying careful attention to language use, because the words we use are part of the way we think. For Dennett, as soon as you say that something ‘enters consciousness’, for example, or ‘reaches the threshold of consciousness’ – phrases that the neuroscientific literature on consciousness is full of, once you start to notice them – you are creating a ‘Cartesian Theatre’. You are imagining that being conscious – enjoying that apparently rich and unified feeling of being me now – is like being the audience of the show on the stage of a special mental theatre (a new version of Ryle’s ghost). We will return to this idea in Chapter 5, but for now the important thing to bear in mind is that theories and statements about consciousness may be implicitly or explicitly presented as materialist, yet be something else when you dig a little deeper. Dennett says that ‘accepting dualism is giving up’ (1991, p. 37). But avoiding it is not easy.

Monist theories try to avoid it, some by claiming that the mental world is fundamental, others that the physical world is. So, for example, you might doubt that real pencils actually exist out there and decide that only ideas or perceptions of pencils exist – making you a mentalist or an idealist. This does away with the awkward division but makes it very hard to understand why physical objects seem to have enduring qualities that we can all agree upon – or indeed how science is possible at all. Even so, there have been many philosophical theories of this kind. The British empiricist George Berkeley, for example, replaced matter with sensations in minds.

At the other extreme are materialists who argue that there is only matter (or physicalists, who include energy as well as matter), and that the physical universe is causally closed. This means that the laws governing the interactions between matter and energy exhaust all the forces of the universe, so there is no room for non-physical minds or consciousness to intervene. Materialism includes identity theory (which makes mental states identical with brain states) and functionalism (which equates mental states with functional states). In these theories there is no mind, or mental force, apart from matter.

Some people find materialism unattractive as a theory of consciousness because it seems to take away the very phenomenon, subjective experience, that it was trying to explain. In particular, the powerful feeling we have that our conscious decisions *cause* our actions is reduced to purely physical cause and effect. Another problem is the difficulty of understanding how thoughts and feelings and mental images can really *be* matter when they seem to be so different. Materialism makes it hard to find any way of talking about consciousness that does justice to the way it *feels*.

However, materialism does not necessarily imply that consciousness can be *reduced* to physical properties. For example, consciousness might not be identical with physical properties, but nonetheless depend on nothing other than physical properties – that is, *supervene* on physical properties. This means that there can be no mental difference without some physical difference: any difference in consciousness must be accompanied by a difference in the brain, but the reverse is not true. So, the same conscious experience might be possible given two different brain states. But although supervenience may help us avoid some of the problems of materialism, it leaves unspecified the precise way in which consciousness depends on the physical (Francescotti, 2016): is the dependence logical, causal, constitutive, or in fact a matter of genuine identity?

The doctrine of ‘epiphenomenalism’ is the idea that mental states are produced by physical events but have no causal role to play. In other words, physical events cause or give rise to mental events, but mental events have no effect on physical events. This idea is sometimes attributed to Julien Offray de La Mettrie (1748), whose book *L’homme machine (Man a Machine)* horrified eighteenth-century French readers. He claimed that like those of other animals, human bodies are clever machines and that ‘the soul’s various states are always correlated with the body’s’. He called this correlation a dependence, and one whose causes and effects our ‘feeble understanding’ was not yet able to unravel (p. 8). But later he also described the mind–body connection in terms of identity, placing him somewhere between epiphenomenalism and materialism: ‘since all the soul’s abilities depend so much on the specific organisation of the brain and of the whole body that obviously they *are* nothing but that very organisation, the machine is perfectly explained!’ (p. 22). Thomas Henry Huxley, the English biologist and palaeontologist who did so much to promote Darwin’s theory of evolution by natural selection, was one of the best-known epiphenomenalists. He did not deny the existence of consciousness or of subjective experiences, but he denied them any causal influence. They were powerless to affect the machinery of the human brain and body, just as the sound of a locomotive’s steam-whistle cannot influence its machinery, and a shadow cannot affect the person who casts it. He referred to animals, including humans, as ‘conscious automata’.

• SECTION ONE: THE PROBLEM

One problem with epiphenomenalism is this: if conscious experiences can have no effect on anything whatsoever, then we should never know about or be able to speak about them, since this would mean they had had an effect. Another difficulty is that if mind is a by-product or side-effect of the physical world but is not actually physical itself, then epiphenomenalism is really a kind of dualism. Nevertheless, scientific or methodological behaviourism is built on one version of this idea: the idea that mental states exist, but do not have effects that can be (or need to be) investigated scientifically.

Trying to avoid the extremes of materialism and idealism without falling into dualism are various kinds of 'neutral monism', which claim that the world is all made of one kind of stuff, but a stuff that cannot be classified as either mental or physical. William James started with 'the supposition that there is only one primal stuff or material in the world, a stuff of which everything is composed' (1904, p. 477). To avoid reducing mind to matter or doing away with matter altogether, he suggested that instead of thinking of a world of physical objects, we should think of a world of possible and actual sense-data, in which the present is made of 'pure experience', before consciousness and content get retrospectively split off from each other. 'A science of the relations of mind and brain must show how the elementary ingredients of the former correspond to the elementary functions of the latter' (1890, i, p. 28), he said, but he did not underestimate the difficulty of this task. The difficulty of developing a detailed account of the neutral *stuff* which the theory depends on, and the fact that it 'attracts neither those who think the mental is a basic feature of reality, nor those who dream of the desert landscape of physics' (Ludwig, 2002, p. 21), together make it a generally unpopular view.

Another way of trying to get around the problem is panpsychism, the view that all material things have awareness or mental properties, however primitive. If materialism is the thesis, Chalmers says, and dualism is the antithesis, panpsychism is the synthesis (Chalmers, 2017). In some versions everything in the universe is conscious, including electrons, clouds, rivers, and cockroaches. This 'pure' panpsychism (Strawson, 2006, 2008) can be thought of as like carrying out a 'global replace' on the usual definition of physical things (mass and energy) as being non-experiential and defining them as being experiential as well (Strawson, 2011, p. 271). This leaves in place everything currently explained by physics. In other versions, experience is another fundamental quality, alongside matter and energy, to be added to our understanding of the world.

Panpsychism raises difficult questions. Is a stone aware? If so, is each of its molecules also separately aware? Are the loose bits on the edge of the stone separately aware when they are just hanging on or only when they are completely knocked off? What would it mean for something as simple as an electron to have mental attributes? Despite these difficulties, some popular recent theories of consciousness, including Integrated Information Theory (Chapter 5), are considered to be forms of panpsychism (Tononi and Koch, 2015).

Given the difficulty of uniting the world, it is not surprising that dualism remains enduringly popular despite its problems. Given the difficulties that arise as soon as we even try to talk about mind and matter, it is also unsurprising that the whole field of psychology has had such trouble with the concept of consciousness.

CONSCIOUSNESS IN PSYCHOLOGY

The term 'psychology' first appeared in the eighteenth century to describe the philosophy of mental life, but it was towards the end of the nineteenth century that psychology first became a science, distinguished from philosophy by being based primarily on empirical data. At that time several different approaches to the study of the mind were emerging. Some were more concerned with physiology and the idea of psychology as an objective science, and some were more concerned with studying subjective experience, but there was, as yet, no great split between the two.

William James's classic text *The Principles of Psychology* (perhaps the most famous book in the history of psychology), begins 'Psychology is the Science of Mental Life, both of its phenomena and their conditions' (1890, i, p. 1). James includes among these phenomena feelings, desires, cognitions, reasonings, and volitions – in other words, the stuff of consciousness. Another textbook from James's era defines psychology, or 'Mental Science', as

the science that investigates and explains the phenomena of mind, or the inner world of our conscious experience. These phenomena include our feelings of joy and sorrow, love, etc., [. . .] our conscious impulses and volitions, our perceptions of external objects as *mental* acts, and so forth.

(Sully, 1892, i, p. 1)

With his monist approach, James dismissed the dualist concepts of soul or 'mind-stuff', and quickly pointed out that consciousness can be abolished by injury to the brain and altered 'by a very few ounces of alcohol or grains of opium or hasheesh' (1890, i, p. 4). So, he assumed that a certain amount of brain physiology must be included in psychology. Nevertheless, consciousness was at the heart of his psychology. He popularised the phrase 'the stream of consciousness' (which was perhaps first used by the English philosopher Shadworth Hodgson in 1865; see Billig, 2012) to describe the apparently ever-changing and flowing succession of thoughts, ideas, images, and feelings. His psychology was therefore very much an integrated science of mental life. Consciousness was at its heart, but was not divorced either from the results of experiments on attention, memory, and sensation, or from physiological study of the brain and nervous system. His brother, the novelist Henry James, experimented with the 'stream-of-consciousness' style of writing that became an important part of Modernist literature, giving readers access to places and events and characters only through the filter of a central character's consciousness.

After he had gone she leaned back in her chair and closed her eyes; and for a long time, far into the night and still further, she sat in the



ACTIVITY 1.1

Defining consciousness

There is no generally recognised definition of consciousness, which is why we have not given one here. See whether you can create your own.

First, get into pairs. One person proposes a definition of consciousness. Then the other finds something wrong with it. Don't be shy or think too long – even the silliest suggestions can be fun to try. So just throw up one idea and wait for it to be knocked down. Then swap over. Do this as quickly as you reasonably can until each of you has had several turns.


Get back together into the group and find out what kinds of objections you all came up with.

Why is defining consciousness so hard when we all think we know what it is?



PROFILE 1.2

William James (1842–1910)



William James was born in New York, the eldest of five children, one of whom was the novelist Henry James. When they were young their wealthy father took the family travelling all over Europe, educating them intermittently along the way. James continued his transatlantic travels most of his life, speaking several languages fluently and getting to know the foremost scholars and scientists of his day. At 18 he wanted to be a painter, and then after long bouts of despair and depression he studied medicine at Harvard, where he eventually taught physiology, psychology, and philosophy. He married in 1878 and was a devoted family man. His *Principles of Psychology* (1890), based on twelve years of introspective investigations, has been called ‘the best-known book in all psychology’ (Gregory, 1986, p. 395). It made him famous for such phrases as ‘the stream of consciousness’ and ‘the specious present’. His other books include *The Varieties of Religious Experience* (1902) and *Pragmatism* (1907). He was a firm believer in free will and a personal spiritual force. He died of heart disease at his summer home in New Hampshire.

still drawing-room, given up to her meditation. [...] she had seen only half his nature then, as one saw the disk of the moon when it was partly masked by the shadow of the earth. She saw the full moon now – she saw the whole man. [...] she lingered in the soundless saloon long after the fire had gone out. There was no danger of her feeling the cold; she was in a fever.

(Henry James, *The Portrait of a Lady*, 1881)

William James was able to build on a large body of research in anatomy, physiology, and psychophysics. Psychophysics is the study of the relationship between physical stimuli and reportable sensations – or, you could say, between events and experience. Psychophysicists such as Ernst Weber and Gustav Fechner studied the relationships between physical luminance and perceived brightness, weight and sensations of heaviness, and sound pressure and loudness. From this research came the famous Weber–Fechner Law relating sensation to the intensity of a stimulus. Fechner also wanted to be able to relate sensations to excitations within the brain, but in his time this was simply not possible.

If James’s *Principles* helped found the modern science of psychology in North America, the experimental work being carried out in Germany was creating a similar movement on the other side of the Atlantic. In 1850 Hermann von Helmholtz made the first measurement of the conduction speed of nerve signals. This was popularly referred to as the ‘velocity of thought’, although in fact he had measured peripheral processes and reaction times, and argued that conscious thought and the interaction of physical and mental processes goes on in the brain. He was especially interested in visual illusions and

the tricks that our senses can play, and he proposed the novel idea that what we see is not the representation of the stimulus on the retina but is determined by inferences and expectations.

The empiricism of the likes of Helmholtz set the scene for another important trend in the European history of the psychology of consciousness: phenomenology. Phenomenology cuts across our neat philosophy/psychology divide, because it is both a philosophy and a psychology based on putting subjective experience first. The German philosopher Edmund Husserl argued for going back to ‘the things themselves’. By this, he meant that we should go back to the ways that things are actually given in experience, not back to the physical objects which the empiricist would work with. The method he proposed was a systematic inquiry into immediate conscious experience. This inquiry was to be done without preconceptions, by suspending or ‘bracketing’ any scientific and logical inferences about the world. This suspension of judgement he called the phenomenological reduction or epoché (see Chapter 17).

Husserl's phenomenology built on the earlier work of Franz Brentano, whose theory of consciousness was based on the idea that every subjective experience is an act of reference. Conscious experiences are *about* objects or events, while physical objects are not about anything. For example, I might have a belief about horses, but a horse itself is not about anything. This 'aboutness' he called 'intentionality'.

It is most important to realise that this awkward word gets used in many different senses. By and large philosophers use it in Brentano's sense, as meaning reference or aboutness. In psychology (and in ordinary language when it is used at all), intentionality usually means 'having intentions' or having plans or goals or aims. If you come across this word, ask yourself which meaning is intended, so you can avoid getting confused and can spot some of the amusing muddles created by people who mix them up.

The idea that all experience is about something is also questionable in itself. Some claim that it is possible to have 'pure consciousness', consciousness without being conscious of anything (Chapter 18). We might also ask whether all emotions or sensations (joy, heat) are *about* or *referring to* things. And if so, how do they relate to consciousness? Am I conscious of an emotion, so is my consciousness about something which is itself already about something?

A separate approach to studying subjective experience used methods based on introspection, or self-observation, which the German physiologist Wilhelm Wundt helped develop. Wundt had founded the first laboratory of experimental psychology in 1879, and for this he is often called the father of experimental psychology. While the physiology in which he was trained studied living systems from the 'outside', he wanted to build a psychology based on studying from the 'inside' – in other words, introspection. Like Husserl, he insisted that introspective study had to be systematic and rigorous, and so he trained people to make precise and reliable observations of their own experience. Later researchers, such as Wundt's student Edward Titchener, investigated other ways of making use of introspection in science, primarily studying sensation and attention.

Wundt claimed to find that there were two kinds of 'psychical elements': the *sensory* elements or simple sensations such as tones, heat, or light, and the *affective* elements or simple feelings, such as the sensory pleasure or displeasure that might accompany the simple sensations. Every conscious experience depended on a union of these two types. Like many others around this time, he hoped to be able to build up a science of consciousness by understanding the units or atoms which, combined into complex compounds, made up 'the actual contents of psychical experience' (1897, p. 29) – an atomistic approach to consciousness that William James utterly rejected.

Although phenomenology and introspectionism both had the benefit of dealing directly with experience (or at least, with what people said about their experience), they faced serious difficulties. For example, Wundt's trained participants had to look at a colour or listen to a ticking metronome and report their thoughts and feelings, but reporting can itself interfere with thoughts and feelings, and some of them might not have described their feelings accurately or truthfully – and with no objective measure it was impossible to find out. These were among

• SECTION ONE: THE PROBLEM

the reasons why introspectionism fell out of favour and behaviourism became so successful, though less so in Europe than in the United States.

The founder of behaviourism, American psychologist John B. Watson, wrote in 1913: 'Psychology, as the behaviorist views it, is a purely objective, experimental branch of natural science which needs introspection as little as do the sciences of chemistry and physics' (p. 158). He proposed to abolish such nonsense as introspection and consciousness, and establish a psychology whose goal was the prediction and control of behaviour. One advantage of this new approach was that behaviour can be measured much more reliably than introspections can. Also, human psychology could build on the considerable knowledge of the behaviour of other animals. As Watson proclaimed, behaviourism 'recognizes no dividing line between man and brute' (p. 158).

Although Watson is usually credited with – or blamed for – the expulsion of consciousness from psychology, similar views were already gaining ground long before. In 1890 James wrote: 'I have heard a most intelligent biologist say: "It is high time for scientific men to protest against the recognition of any such thing as consciousness in a scientific investigation"' (1890, i, p. 134). Watson also exaggerated the dominance of 'introspectionism' as a scientific movement, as well as the naïvety of Wundt's understanding of introspective methods, to make his own 'revolution' seem more dramatic (Costall, 2006).

Watson built many of his ideas on the work of Ivan Pavlov, the Russian physiologist famous for his work on reflexes and classical conditioning. He studied the way that repetition increased the probability of various behaviours and assumed that almost everything we do, including language and speech, is learned in this way. Subsequently, the emphasis in behaviourism shifted to the study of operant conditioning, with B.F. Skinner's studies of rats and pigeons that learned by being rewarded or punished for their actions. For Skinner,



FIGURE 1.4 • When the rat presses the lever it may receive a food pellet or a sip of water. Rats, pigeons, and many other animals can easily learn to press a certain number of times, or only when a green light is on, or when a bell sounds. This is known as operant conditioning. Some behaviourists believed that studying animal learning was the best way to understand the human mind.

human behaviour was shaped by the history of reinforcements, and he believed that with the right reinforcement schedules a human utopia could be created (Skinner, 1948). As for consciousness, he believed it was just an epiphenomenon and its study should not be the task of psychology. In the words of Watson's biographer David Cohen, 'Behaviourism was a self-conscious revolution against consciousness' (1987, p. 72).

Behaviourism was enormously successful in explaining some kinds of behaviour, particularly in the areas of learning and memory, but it more or less abolished the psychological study of consciousness, and even the use of the word 'consciousness' became unacceptable. Also, though it arguably generated valuable reflections on the nature of evidence and behaviour and objectivity, behaviourism threw out the much more even-handed mind-and-body approach of William James's 'science of mental life'. And although thinkers like Maurice Merleau-Ponty made detailed studies of perception and embodiment that would eventually help bring phenomenology closer to psychology, in the early twentieth century phenomenology was increasingly alienated from a mainstream psychology dominated by behaviourism. All this led to half a century of a very restricted kind of psychology indeed.

By the 1960s, behaviourism was losing its power and influence, and cognitive psychology, with its emphasis on internal representations and information processing, was taking over, but 'consciousness' was still something of a dirty word. In his widely read history *Psychology: The Science of Mental Life*, George Miller warned:

Consciousness is a word worn smooth by a million tongues. Depending upon the figure of speech chosen it is a state of being, a substance, a process, a place, an epiphenomenon, an emergent aspect of matter, or the only true reality. Maybe we should ban the word for a decade or two until we can develop more precise terms for the several uses which 'consciousness' now obscures.

(1962, p. 40)

No one got quite as far as banning its use, but it was certainly more than a decade before the word 'consciousness' became acceptable again in psychology. The change was due partly to growing interest in big questions about experience beyond the everyday and the individual: questions about spiritual experience, about drug-induced states, about mental illness, hypnosis, and the paranormal. One route these interests took was via William James's 1902 book *The Varieties of Religious Experience*, which later inspired other books like *The Varieties of Psychedelic Experience* (Masters and Houston, 1967) and *The Varieties of Scientific Experience* (Sagan, 2006). In the course of his career James gradually developed a new form of philosophy, called radical empiricism, which insisted that experience must always be at the heart of philosophical inquiry, and that experience has to be understood as fundamentally about meaning, not just physical data. The work of James, Carl Jung, and others contributed to the explicit focus on spirituality and transcendence in transpersonal psychology, and this, together with the rise of the counter-cultures of the 1960s, created other paths for consciousness to creep back into the academy. During the 1970s, research on mental imagery (see Chapter 5) and altered states of consciousness such as sleep and drug-induced states (Section Five), and the beginnings of computer science (Chapter 12),

'Maybe we should ban the word for a decade or two'

(Miller, 1962, p. 40)

'genteel avoidance of consciousness [...] feels much like tiptoeing to keep from waking the insane attic-bound Aunt of a Gothic novel'

(Banks, 1993, p. 257)

• SECTION ONE : THE PROBLEM

opened things up further. But it was nearly three decades before the sudden explosion of interest in consciousness in the 1990s.

From around the 1950s to the 1990s the 'first-generation' cognitive sciences had conceived of the mind in terms of abstract, language-like representations (Lakoff and Johnson, 1999, pp. 77–78) and relied heavily on analogies with digital computers, but increasingly people began to think more in terms of interconnected networks that change over time. From this connectionist approach came the idea of the neural network, which revolutionised the study and creation of artificial intelligence (see Chapter 12). This movement, along with the embodied philosophy of Merleau-Ponty, contributed to the emergence of the 'second generation' of cognitive science, which recognised that brains are always found in bodies, and bodies in environments – both physical and social.

And thinking about cognition as '4E' (Menary, 2010) – as embodied, enactive, embedded, and extended (involving other objects and people in the environment) – opens up much more space for experience than does a computational brand of cognitivism. As the authors of *The Embodied Mind* put it: in the embodied paradigm, 'cognition and consciousness – especially self-consciousness – belong together in the same domain. Cognitivism runs directly counter to this conviction [. . .] for cognitivists, cognition and intentionality (representation) are the inseparable pair, not cognition and consciousness' (Varela et al., 1991, p. 173). Thinking about the kinds of experience that come from having a body with particular sensory and motor capacities, and from the feedback between these capacities and the environment, gives us an alternative to trying to uncover consciousness through the neurons alone.

'The mind is not in the head'

(Varela, 1999, p. 72)

An approach that combines the insistence on feedback between brain, body, and world with a firm basis in brain function is that of predictive processing, the idea that brains are essentially prediction machines, constantly trying to match incoming sensory inputs with their own expectations or predictions (Clark, 2013). This is a modern version of Helmholtz's idea of 'unconscious inference' and British psychologist Richard Gregory's (1966/1997) much later notion that perceptions are guesses, or hypotheses, about the world. The difference is that with advances in neuroscience and computation we can begin to work out how the embodied brain builds its predictions and adapts the body's responses to the world.

Dynamics-based paradigms like this give us challenging new ways of thinking about the status of the brain in relation to the mind and consciousness, which we will return to in Chapter 3. In principle, they also link to other theories that stress the contexts of consciousness, such as social constructionism, a movement which built on the developmental psychology of the Soviet psychologist Lev Vygotsky in the 1930s, and which investigates how reality as we know it is constructed through social interactions. They have certainly begun to generate resonances within cognitive approaches to the study of literature (Caracciolo and Kukkonen, 2014). But in practice, the methods and questions of 4E and predictive processing have not yet tended to expand very far into the study of language, history, and culture.

Even now, after centuries of philosophical and psychological inquiry, our understanding of how behaviour and 'introspection' relate to consciousness leaves a lot to be desired (Costall, 2006), and even a definition of consciousness that everyone

can agree on remains out of reach (Dietrich, 2007). But at least we are now allowed to talk about it.

In this book we use 'consciousness' to mean subjective experience. We use the word 'awareness' to mean the same thing, and will often also use the phrase 'what it's like' to get at how it feels to you (see Chapter 2). What we are trying to understand is the nature and origins of your experience of that pencil, or anything else.

'Human consciousness is just about the last surviving mystery'

(Dennett, 1991, p. 21)

THE MYSTERIOUS GAP

'Human consciousness is just about the last surviving mystery,' says Dennett (1991, p. 21). He defines a mystery as a phenomenon that people don't know how to think about – yet. Once upon a time the origin of the universe, the nature of life, the source of design in the universe, and the nature of space and time were all mysteries. Now, although we do not have answers to all the questions about these phenomena, we do know how to think about them and where to look for answers. With consciousness, however, we are still in that dreadful – or delightful – state of mystification. Our understanding of consciousness is a muddle.

The cause of that mystification, as we have seen in our quick look at the history of consciousness, seems to be a gap. But what sort of a gap is it?

"'A motion became a feeling!' – no phrase that our lips can frame is so devoid of apprehensible meaning.' This is how William James describes what he calls the "'chasm" between the inner and the outer worlds' (1890, i, p. 146). Before him, Tyndall had famously proclaimed: 'The passage from the physics of the brain to the corresponding facts of consciousness is unthinkable' (James, 1890, i, p. 147). In *The Nervous System and the Mind*, Charles Mercier referred to 'the fathomless abyss that separates mind from matter', but also advised the student of



FIGURE 1.5 • The fathomless abyss

• SECTION ONE: THE PROBLEM

psychology to ponder the fact that a change of consciousness never takes place without a change in the brain, and a change in the brain never without a change in consciousness.

Having firmly and tenaciously grasped these two notions, of the absolute separateness of mind and matter, and of the invariable concomitance of a mental change with a bodily change, the student will enter on the study of psychology with half his difficulties surmounted.

(Mercier, 1888, p. 11)

'Half his difficulties ignored, I should prefer to say,' remarks James. 'For this "concomitance" in the midst of "absolute separateness" is an utterly irrational notion' (1890, i, p. 136). He quotes the British philosopher Herbert Spencer as saying,

Suppose it to have become quite clear that a shock in consciousness and a molecular motion are the subjective and objective faces of the same thing; we continue utterly incapable of uniting the two, so as to conceive that reality of which they are the opposite faces.

(1890, i, p. 147)

To James it was inconceivable that consciousness should have nothing to do with events that it always accompanied. He urged his readers to reject both the epiphenomenalist/materialist automaton theory and the dualist 'mind-stuff' theory and, in the terms of his neutral monism, to ponder the how and why of the relationship between physiology and consciousness (James, 1904).

As we have seen, the automaton theory gained ground, and behaviourism, with its thorough-going rejection of consciousness, held sway over most of psychology for half a century or more. Behaviourists had no need to worry about the great gulf because they simply avoided mentioning consciousness, subjective experience, and inner worlds. It was only when this period drew to a close that the problem became obvious again. In 1983 the American philosopher Joseph Levine coined the phrase 'the explanatory gap,' describing it as 'a metaphysical gap between physical phenomena and conscious experience' (Levine, 2001, p. 78). No sooner had consciousness been allowed back into science than the mysterious gap had opened up once more.

'The hard problem [. . .] is the question of how physical processes in the brain give rise to subjective experience'

(Chalmers, 1995b, p. 63)

Then in 1994 a young philosopher, David Chalmers, presented a paper at the first Toward a Science of Consciousness conference in Tucson, Arizona. Before getting into the technicalities of his argument against reductionism, he wanted to clarify what he thought was an obvious point: that the many problems of consciousness can be divided into the 'easy' problems and the truly 'hard problem'. To his surprise, his term 'the hard problem' stuck, soon provoking numerous debates and four special issues in the newly established *Journal of Consciousness Studies* (Shear, 1997).

According to Chalmers, the easy problems are those that are susceptible to the standard methods of cognitive science, and might be solved, for example, by understanding the computational or neural mechanisms involved. They include the mechanisms of attention, behavioural control, and the sleep-wake cycle.

Phenomena like these are in some way associated with the notion of consciousness, but they are not deeply mysterious. In principle (even though it may not really be 'easy') we know how to set about answering them scientifically. The really hard problem, by contrast, is *experience*: what it is like to *be* an organism, or to *be in* a given mental state, to experience the quality of deep blue or the sensation of middle C. 'If any problem qualifies as *the* problem of consciousness,' says Chalmers,

it is this one. [...] [E]ven when we have explained the performance of all the cognitive and behavioral functions in the vicinity of experience – perceptual discrimination, categorization, internal access, verbal report – there may still remain a further unanswered question: *Why is the performance of these functions accompanied by experience?* [...] Why doesn't all this information-processing go on 'in the dark', free of any inner feel? In other words, 'Why should physical processing give rise to a rich inner life at all?'

(1995a, pp. 201–203)

Stated at its most succinct: 'The hard problem [...] is the question of how physical processes in the brain give rise to subjective experience' (Chalmers, 1995b, p. 63). Or, as British philosopher Colin McGinn puts it: 'How can technicolour phenomenology arise from soggy grey matter?' (1991, p. 1). This is the latest incarnation of the mysterious gap.

CONSCIOUSNESS IN CONTEXT

One of the reasons why the mysteries of consciousness are so hard and also so enticing to grapple with is that they are so closely linked to what it means to be me: asking 'am I conscious now?' or 'what is it like to be me now?' leads naturally on to the questions 'what am I?', 'who is asking the question?', and 'what am I doing?' (Blackmore, 2011), and once we tackle these, we find ourselves confronting the concepts of self and free will. We will address the problem of free will head-on in Chapter 9, once we have explored how the mechanisms of attention and embodied action contribute to our sense of agency in the world. The self will pop up in all sorts of contexts as we go along, but we will delay a thorough investigation of it until the final section, where we will bring together evidence from the many different fields to which it relates, and ask what its uses and its pitfalls are as a concept.

'How can technicolour phenomenology arise from soggy grey matter?'

(McGinn, 1991, p. 1)



CONCEPT 1 THE HARD PROBLEM

- The hard problem is to explain how physical processes in the brain give rise to subjective experience. The term was coined in 1994 by David Chalmers, who distinguished it from the 'easy problems' of consciousness. These include the ability to discriminate, categorise, and react to stimuli, or to report mental states, focus attention, or control behaviour; the integration of information by cognitive systems; and the difference between wakefulness and sleep. By contrast, the hard problem concerns experience itself – that is, *subjectivity* or 'what it is like to be. . . .'

The hard problem can be seen as a modern version, or aspect, of the traditional mind–body problem. It is the problem of how to cross the 'fathomless abyss' or 'chasm', or how to bridge the 'explanatory gap' between

the objective material brain and the subjective world of experience.

Some argue that new physical principles are needed to solve the hard problem. Mysterians believe it can never be solved; illusionists think that, like 'consciousness itself', it is illusory; and many neuroscientists believe that once we solve the easy problems, the hard problem will disappear.

Underlying consciousness, the self, and free will is a notion that seems the dark flip-side of them all: the unconscious. The history of the unconscious has been a stormy one. The idea that much of what goes on in the nervous system is unconscious and that our conscious experiences depend upon unconscious processing seems quite natural to us today. Yet it was deeply disturbing to many nineteenth-century scientists, who assumed that inference and thinking, as well as ethics and morality, require consciousness. To them, the idea that thinking could happen without consciousness seemed to undermine the moral or spiritual superiority of 'Man'. This meant that the notion of the unconscious which was derived from physiological studies of the time, such as Helmholtz's idea that perceptions are 'unconscious inferences' and James's (1902) talk of 'unconscious cerebration', was genuinely shocking.

The notion of the unconscious developed by Sigmund Freud was a crucial part of his 'psychodynamic' theory of how conscious and unconscious forces interact to produce personality and motivation. In Freud's theory the unconscious (in his early work also called the subconscious) consists of the impulses of the 'id', including biological desires and needs; the defence mechanisms and neurotic processes of the 'ego'; and all the mass of unwanted or unacceptable material that is repressed by the 'super-ego' – a part of the mind acquired through education in childhood, and the source of conscience and guilt. All these unconscious feelings, images, and forbidden wishes or instincts might then appear in dreams or cause neurotic symptoms (e.g. Freud, 1915, 1923/1927). Although Freud was trained as a neurologist, and frequently referred to his work as a 'new science', his theories were derived almost entirely from case studies of psychiatric patients and from his own self-analysis, and were largely unfalsifiable. The theories of psychoanalysis have not stood the test of time, and the ethics of Freud's interactions with his patients were dubious, especially when it came to 'recovering' their 'memories' of childhood sexual abuse. Nonetheless, his work did manage to lastingly influence everyday notions of what the unconscious is and does.

That night he had a terrible dream [...]. Fear was the beginning, fear and desire and a horrified curiosity at what was to come. It was night, and his senses were alert, because from far away a turmoil, a roar, a blend of noise approached [...]. But he knew a phrase, dark, yet denoting what was coming: 'The foreign god!' [...] And in the splintered light, from woody hills, between trunks and mossy boulders it rolled and crashed earthward like a vortex: men, animals, a swarm, a raging mob, and flooded the slopes with bodies, flames, tumult, and a delirious dance. [...] Great was his abhorrence, great his fear, honourable his will, to protect to the last what was his from the foreign, from the enemy of the sober and dignified mind. But the din, the howling, multiplied by the echoing cliff face, grew, gained the upper hand, swelled to a ravishing madness. [...] His heart thudded with the drumbeats, his brain gyrated, anger gripped him, blindness, deadening lust, and his soul craved to join the god's dance. The obscene symbol, enormous, wooden, was uncovered and raised: then more riotously they howled the watchword.

(Thomas Mann, *Death in Venice* [*Der Tod in Venedig*], 1912, our translation)

In the late twentieth century, Freud's unconscious was largely replaced by the idea of a 'cognitive unconscious' (Kihlstrom, 1987) capable of subliminal perception and many types of thinking, learning, and memory without awareness, and then by what is sometimes called the 'new unconscious', which expands this notion to emphasise emotions, motivation, and control (Hassin et al., 2005).

We will see later how difficult it is to think about the unconscious without assuming a 'magic difference' between things that are 'in' or 'out' of consciousness, things that have or have not 'reached consciousness' or been made 'available to consciousness'. A wide range of evidence, which we will discuss particularly in Chapters 4 and 8, suggests that we should reject any firm distinction of this kind, but ordinary intuitions about consciousness depend utterly on such a distinction. This is a familiar situation, and just one more reason why the problem of consciousness is so perplexing. The idea of the magic difference will be one of the threads we will need to hold on to if we are to find our way through the maze of theories and intuitions that promise to help us understand consciousness.

Bayne, T., Cleeremans, A., and Wilken, P. (2009). *The Oxford companion to consciousness*. Oxford: Oxford University Press.

Hundreds of short entries by over 200 authors on everything from access consciousness to zombies; provides an idea of the scope of consciousness studies.

Chalmers, D.J. (1995b). The puzzle of conscious experience. *Scientific American*, December, 62–68.


The easiest version of Chalmers's 'hard problem'. For more detail, read Chalmers, 1995a and 1996.

Dennett, D.C. (1991). Explaining consciousness. *Consciousness explained* (pp. 21–42). Boston, MA: Little, Brown.

The mystery of consciousness and the problems of dualism.

Gregory, R.L. (2004). *The Oxford companion to the mind*. Oxford: Oxford University Press.

Short entries on most authors and ideas presented here, and a multi-authored section on consciousness. Non-philosophers will find it helpful for looking up philosophical concepts.



What is it like to be . . . ?

CHAPTER

TWO

BEING A . . .

What is it like to be a bat? This is one of the most famous questions ever asked in the history of consciousness studies. It came to prominence in a 1974 paper of the same name by the American philosopher Thomas Nagel. He argued that understanding how mental states can *be* neurons firing inside the brain is a problem quite unlike understanding how water can be H₂O, or how genes can be DNA. 'Consciousness is what makes the mind-body problem really intractable', he said (1974, p. 435; 1979, p. 165), and by consciousness he meant *subjectivity*. To make this clear he asked, 'what is it like to be a bat?'

Do you think that your cat is conscious? Or the birds outside in the street? Perhaps you believe that horses are conscious but not worms, or living creatures but not stones. We shall return to these questions (in Section Four), but here let us consider what it means to say that another creature is conscious. If you say that the stone is not conscious you probably mean that it has no experiences and no point of view; that there is nothing it is like to *be* the stone. If you believe that your neighbour's new kitten, or the woodlouse you narrowly avoided crunching underfoot, is conscious, then you probably mean that they do have a point of view; there is something it is like to be them.

As Nagel put it, when we say that another organism is conscious we mean that 'there is something it is like to *be* that organism [. . .] something it is like *for* the organism' (1974, p. 436); 'the essence of the belief that bats have experience is that there is something that it is like to be a bat' (1974, p. 438). There is currently no agreement over how to define consciousness (Dietrich, 2007; Nunn, 2009;

Vimal, 2009), so this is probably the closest we can come to a definition – that consciousness is subjectivity, or ‘what it is like to be . . .’.

Here we must be careful with the phrase ‘what it is like. . .’. Unfortunately, there are at least two things we might mean when we ask what something is like in English. Consider the statement ‘this ice cream tastes like rubber’, or ‘his look cut through her like a knife’. In this case we are comparing things, making analogies, or saying what they resemble. This is *not* what Nagel meant. The other meaning is about identity, not comparison, and is found in such questions as ‘What is it like to work at McDonald’s? What is it like to be able to improvise fugues at the keyboard? . . . to be someone inconceivably more intelligent than yourself? . . . to be a molecule, a microbe, a mosquito, an ant, or an ant colony?’ (see Hofstadter and Dennett, 1981, pp. 404–405 for many more such provocative questions). British social psychologist Guy Saunders prefers the less ambiguous phrases ‘How it is to be . . .’ and ‘How it is for you’ (2014, p. 146). In the more commonly used wording, remember that what we are getting at is: what is it like ‘from the inside’?

Now, imagine being a bat. Bats’ experiences must be very different from ours, which is why Nagel chose the bat for his famous question. Their brains, way of life, and sensory systems are well understood (Dawkins, 1986; Akins, 1993). Most use either sound or ultrasound for echolocation, detecting objects by emitting rapid high-pitched clicks which bounce off any objects in the vicinity and then measuring the time taken for the echo to return. Natural selection has found ingenious solutions to the many interesting problems posed by echolocation. Some bats cruise around emitting clicks quite slowly so as not to waste energy, but then when they are homing in on prey or approaching a potential danger the clicks speed up. Many have mechanisms that protect their ears from the loud blast of each click and then open them to receive the faint echo. Some use the Doppler shift (think of the changing frequency of a passing siren) to work out their speed relative to prey or other objects. Others sort out the mixed-up echoes from different objects by emitting downward-swooping sounds. The echoes from distant objects take longer to come back and therefore sound higher than the echoes from nearer objects. In this way we can imagine that a whole bat world is built up in which higher sounds mean distant objects and lower sounds mean nearer objects.

What would this be like? According to Oxford biologist Richard Dawkins (1986), it might be like seeing is for us. We humans do not typically know, or care, that colour is related to wavelength or that motion detection is carried out in the visual cortex. We just see the objects out there in depth and colour. Similarly, a bat would just perceive the objects out there in depth, and perhaps even in some batty, sonar version of colour. Living in this constructed world would be what it is like to be a bat.



FIGURE 2.1 • The leaf-nosed bat uses sonar to navigate, sending out brief pulses of sound and analysing the returning echoes so as to avoid obstacles, detect fruit and other food, and find its mate. What is it like to be this bat?

• S E C T I O N O N E : T H E P R O B L E M

But can we ever know what it would *really* be like for a bat? As Nagel pointed out, the question is not answered by trying to imagine that *you* are a bat. This will not do. It is no good hanging upside down in a darkened room, making little clicks with your tongue, and flapping your arms like wings. Perhaps, if you could magically be transformed into a bat, you would know and could report back? But no, even this won't do. For if *you* were a bat, the bat in question would not be an ordinary bat – what with having your memories and your interest in consciousness. And if you stopped being yourself and became an ordinary bat, then this bat would have no understanding of English and no ability to ask questions about consciousness, and could not tell anyone else what it was like, even though it might know for itself. So, we humans cannot know what it is like to be a bat, even if we believe that there *is* something it is like.

Nagel's question clarifies the central meaning of the term 'consciousness'. This is what the American philosopher Ned Block calls 'phenomenal consciousness', P-consciousness, or phenomenality. He explains that 'Phenomenal consciousness is experience; what makes a state phenomenally conscious is that there is something "it is like" to be in that state'. He distinguishes this from 'access consciousness' or A-consciousness, which is 'availability for use in reasoning and rationally guiding speech and action' (1995, p. 227). Block asks 'whether phenomenal consciousness includes the cognitive accessibility underlying reportability' (2007, p. 481). In other words, is the ability to say something about our experience inherent to conscious experience, or can the experience and the access be separated out?

At first sight this distinction may seem unnecessary, for surely what we are trying to understand is phenomenality, not access. Yet whenever we study phenomenality, we have to listen to what people say, or in other ways use their reports of conscious experiences. It has even been suggested that reportability should be part of our definition of consciousness: 'it might be a good idea, chiefly for pragmatic reasons, if the default meaning of "consciousness" were to become something like "reportable mental content"' (Nunn, 2009, pp. 7–8). This places enormous significance on the role of language and the communicative context in which language is used. But it also leads us to ask whether these accessible and hence reportable 'contents' are all there is to consciousness, or whether we are missing something crucial when we rely on this kind of testimony.

The intuition that there is more to phenomenal experience than can be accessed and reported is easy to evoke. Look around you now and soak up the colours, feelings, and sounds around you. Now try to describe them to yourself. You may get the distinct feeling that there is a lot more in your conscious experience than you can ever describe, or that the process of trying to capture your experiences destroys them. You may feel that something gets lost whether you speak aloud to others about your experiences, talk to yourself about them, or even only think about them. One of the subsections of Block's 'access consciousness' category is 'reflective consciousness' – that is, higher-order reflection about consciousness, or thinking about thinking. Any kind of 'access', whether fully verbalised or not, can leave us with the impression that we are only scratching the surface, or betraying the reality as soon as we try to pin it down.

*as though the abundance of the soul did not sometimes overflow
in the emptiest metaphors, since no one, ever, can give the exact*

measure of their needs, nor of their ideas, nor of their pains, and since human speech is like a cracked cauldron where we beat tunes to make bears dance, when we would like to melt the stars.

(Gustave Flaubert, *Madame Bovary*, 1856, our translation)

Yet even our firmest intuitions can be horribly wrong and lead us astray. So, are there really two kinds of consciousness, or only one? Many theorists reject the distinction (e.g. Baars, 1988; Dennett, 2005; Carruthers, 2015); some say there is only the reportable stuff (e.g. Nunn, 2009); and others agree with Block that there are two distinct kinds (e.g. Alter, 2010; Raffone and Pantani, 2010). We will return to this distinction (Chapters 8 and 17) and attempts to study it experimentally, but for now ‘phenomenal consciousness’ is what this book is all about.

So, what is it like to be you now? Everything we have said so far implies that there is, uncontroversially, something it is like to be you now – that the problems only begin when you start trying to say in words exactly what it is like for you, or asking about what it is like to be someone or something else. But is this right? A thoroughly sceptical approach would mean questioning even that apparently obvious starting point: your own what-it’s-like. We urge you to do this chapter’s ‘practice’ and become a little more familiar with what it is like to be you.



PRACTICE 2.1

WHAT IS IT LIKE TO BE ME NOW?

As many times as you can, every day, ask yourself ‘**What is it like to be me now?**’ If you practised the previous exercise, ‘Am I conscious now?’, you will have got used to remembering the task, and perhaps to opening your mind a little to watch your own awareness.

This new question is important because so many arguments assume that we know, unproblematically, what our own experience is like, that we know our own qualia directly, and that of course we know what it is like to be ‘me’ now. The only way to have an informed opinion on this important point is to look carefully. What is it really like for you, now?

SUBJECTIVITY AND QUALIA

Let us suppose that you are, right now, getting the unmistakable smell of fresh coffee drifting in from the kitchen. The smell may be caused by chemicals entering your nose and reacting with receptors there, but as far as you are concerned the experience is nothing to do with chemicals. It is a . . . well, what is it? You probably cannot describe it even to yourself. It is just how fresh coffee smells. The experience is private, ineffable, and has a quality all its own. These qualities are known, in philosophy, as qualia. The feel of the wind on your cheeks as you

'the smell of spices as you walk past a restaurant, the taste of chocolate, the sensation of jumping into a cold swimming pool or relaxing in a hot bath'

(Andrade, 2012, p. 579)

ride your bike is a quale (pronounced *kwar-lay*). The sight of the bluey pink of the sunset sky is a quale. The indescribable chill of delight you experience every time you hear that minor chord is a quale.

The term was first used in this context by the American philosopher and logician Charles Sanders Peirce in 1866, and then in 1929 was adapted by a student of William James's, Clarence Irving Lewis, who defined qualia as the fundamental building blocks of specifically sensory experience – a slant it retains to this day (Keeley, 2009). The concept of qualia has become mired in confusion, but the basic idea is clear enough. The term comes from the Latin *qualis* (of what kind or qualities) and is used to emphasise quality: to get away from talking about physical properties or descriptions, and to point to experience itself. A quale is what something is *like* (in the sense explained above). Conscious experience can be thought of as consisting of qualia, and then '*The problem of consciousness is identical with the problem of qualia, because conscious states are qualitative states right down to the ground. Take away the qualia and there is nothing there*' (Searle, 1998, p. 21, original italics). The problem of consciousness can thus be rephrased in terms of how qualia relate to the physical world, or how *objective* brains and bodies produce *subjective* qualia. There are many possible ways of constructing an answer to the question posed in this way. The substance dualist believes that qualia (e.g. the smell of coffee) are part of a separate mental world from physical objects (e.g. pots of coffee or brains). The epiphenomenalist believes that qualia exist but have no causal properties. The idealist believes that everything is ultimately qualia. The eliminative materialist denies that qualia exist, and so on.

'[Q]ualia [...] never really existed [...]. There are no atoms, no nuggets of consciousness'

(Metzinger, 2009, pp. 50–51)

You may think it unquestionable that qualia exist. After all, you are right now experiencing smells, sounds, and sights, and these are your own ineffable qualia, aren't they? Many theorists would agree with you, but many others would not. The disagreements come about partly because people define the term in different ways: they may use it to refer (amongst other things) to qualities of experience in general, to qualities of sensory experience in particular, to distinct irreducible nuggets of experience, or to ineffable qualities about which their experiencing subjects cannot be mistaken.

In his essay 'Quining qualia', Daniel Dennett sets out 'to convince people that there are no such properties as qualia' (1988, p. 42), and what he rejects is this last use of the term. Dennett 'throws the qualic baby out with the bath water', says British psychologist Jeffrey Gray (2004, p. 153). However, Dennett does not deny the reality of conscious experience as something that has properties, nor does he deny that we say things and make judgements about our own experiences. What he does deny is the existence of the ineffable, intrinsic, private, directly apprehensible 'raw feels' that he claims people tend to mean when they talk about qualia.

Dennett provides many 'intuition pumps' (his term for thought experiments designed to draw intuitions to the surface) to undermine this natural way of thinking. Here is a simple one. The experienced beer drinker says that beer is an acquired taste. When he first tried beer he hated the taste, but now he has come to love it. But which taste does he now love? No one could love that first taste – it tasted horrible. So, he must love the new taste, but what has changed? If you think that there are two separate things here, the actual quale (the way it *really* tastes to him) and his opinion about the taste, then you must be able to decide which

has changed. But can you? And if you admit that opinions can have an effect on *actual* tastes, those *actual* tastes lose the quality of indivisible self-sufficiency that qualia are traditionally thought to have. We normally think in a confused and incoherent way about how things seem to us, claims Dennett, and the concept of qualia just confuses the issue further. It may be, as many philosophers claim, that it is difficult to deny the existence of qualia, but we should try, because 'contrary to what seems obvious at first blush, there simply are no qualia at all' (1988, p. 74).

One of the most common responses to Dennett's argument is that he has constructed a straw-man version of qualia which no one believed in anyway. This is a constant problem when it comes to qualia: what *is* the version that people believe in? There is little consensus about what the term means or why it is needed at all, so its use may confuse more than it clarifies. Perhaps the qualities we so struggle to put into words are qualities of the things we have experiences of (the wind, the sky, the minor chord) rather than of our experiences themselves (Harman, 1990). Perhaps what qualia really offer is a way of making it philosophically acceptable to talk about 'how it feels'. The trouble is that this may also tempt us into thinking that the impressive-sounding qualia are more special, more mysterious, and more totally separate from physical stuff than is necessarily the case.

So, when you next come across the term *qualia* in an argument about consciousness, look closely at how it is used. Is a definition given, or is its meaning taken for granted? If the word is defined, is it a helpful definition, or more of a paraphrase in terms that would have been perfectly good on their own? And how does the definition help to support, or undermine, the argument that is being made?

Even if we could agree on a precise and workable definition of qualia to make it preferable to plainer alternatives (the subjective experience, the what-it's-like), how could we decide whether qualia really exist or not? We cannot do experiments on qualia, at least not in the simple sense of first catching a quale and then manipulating it in the lab. That is the whole point of qualia, of raw feels and the qualities of experience: they do not have physical properties that can be measured. We can, however, do thought experiments.

Thought experiments are, as the name implies, experiments done with the mind. It is important to be clear about their purpose. In an ordinary experiment you manipulate something in order to get an answer about the world. If you do the experiment properly, you may get a reliable answer that is widely applicable and that helps decide between two rival theories. Thought experiments are designed not to manipulate the world or provide definitive answers, but rather to manipulate minds and clarify thinking.

Einstein famously imagined riding on the back of a light wave, and from this idea came some of his theories about relativity and the speed of light. Most thought experiments are, like that one, impossible to carry out, although some end up turning into real experiments as technology changes. Most philosophical thought



FIGURE 2.2 • Is this an ineffable quale?

experiments are of the impossible kind. They have not been done, cannot be done, will never be done, and do not need to be done. Their function is to make you think.

WHAT IS IT LIKE BEING ME NOW?

One of the best-known of such thought experiments gets right to the heart of the problem of consciousness. Are subjective experiences something separate from the brain? Do they make any difference? Does consciousness contain information above and beyond the neural information and other physical states it depends on? Mary may help.

MARY THE COLOUR SCIENTIST

Mary is a brilliant scientist who lives in the far future. She

specialises in the neurophysiology of colour vision and acquires, let us say, all the physical information there is to obtain about what goes on when we see ripe tomatoes, or the sky, and use terms like ‘red’, ‘blue’, and so on.

(Jackson, 1982, p. 130)

She knows everything there is to know about the mechanics of colour perception, about the optics of the eye, the properties of coloured objects in the world, and the processing of colour information in the visual system. She knows exactly how certain wavelengths of light stimulate the retina and travel up the optic nerve to the lateral geniculate nucleus and then on to primary visual cortex and other visual and related areas, eventually producing the contraction of the vocal cords and expulsion of air that results in someone saying, ‘the sky is blue’. But Mary has been brought up all her life in a black-and-white room, observing the world through a black-and-white television monitor. She has never seen any colours at all.

One day, Mary is let out of her black-and-white room and sees colours for the first time. What happens? Will she gasp with amazement and say, ‘Wow – I never realised red would look like *that!*’, or will she just shrug and say ‘That’s red, that’s green, nothing new, of course’? You may like to think about your own answer, or do the group activity, before reading on.

The philosopher Frank Jackson (1982) devised the Mary thought experiment in support of what is called the ‘knowledge argument’ against physicalism. Confessing to being a ‘qualia freak’, he argued that when she comes out she obviously learns something fundamentally new – what red is like. Now she has colour qualia as well as all the physical facts about colour. As Chalmers puts it, no amount of knowledge about, or reasoning from, the physical facts could have prepared her for the raw feel of what it is like to see a blue sky or green grass. In other words, the physical facts about the world are not all there is to know, and therefore materialism has to be false. This way of telling Mary’s story is illustrated powerfully by David Lodge’s version, in his novel *Thinks...*




FIGURE 2.3 • What does Mary say when she finally emerges from her black-and-white room?

That day had dawned – whatever dawn looked like – and now it was the eleventh hour. The minute-hand twitched forward to point to twelve. The clock began to strike. Even louder it seemed to Mary was the beating of her own heart, always prone to palpitations at moments of strong emotion. She heard the sound of bolts being drawn on the other side of the door. She rose from her seat and clasped her bosom with an involuntary movement of one gloved hand.

(David Lodge, *Thinks. . .*, 2001, p. 157; read on in *Thinks. . .* to find out what Lodge imagines happening to Mary next)

If you think Mary will be surprised, are you forced to reject materialism and adopt dualism? Chalmers does so, but there have been many objections to this conclusion, and many other ways of using the thought experiment. For example, some have argued that Mary comes to know an old fact in a new way or from a new viewpoint, or to connect up old facts in new ways, or that she learns a new skill rather than a new fact (for a philosophical overview, see Chalmers, 1996). This sort of argument allows you to think that Mary really does experience something surprising when she comes out – but not because there are irreducibly *subjective* facts in the world.

An alternative is to deny that Mary will be surprised. The philosopher Christopher Maloney (1985) suggests a simple test. Choose a colour (say a nice shade of pale mauve) and give Mary a detailed neurophysiological description of the state associated with seeing that colour. If Mary really does understand all there is to know about the physical nature of colour vision, she must be perfectly well able to imagine what seeing that particular shade of mauve would be like. Then expose her to a range of colour samples and get her to select the pale mauve she had imagined. Maloney believes that she would pass this test. Paul Churchland proposes a related test: give Mary a stimulus that produces in her the relevant state ‘and see whether she can identify it correctly *on introspective grounds alone*, as “a spiking frequency of 90 hz: the kind a tomato would cause” (1985, p. 26). He too thinks it likely that Mary might pass the test. Along similar lines, Dennett argues that this story is not the good thought experiment it appears to be, but a misleading intuition pump that lulls us into vividly imagining (as Lodge does) the event of Mary’s release from the room, and encourages us to misunderstand its



ACTIVITY 2.1

Mary the colour scientist

When Mary comes out of the black-and-white room, will she learn anything new? Will she be surprised at what colours are *like*? Or does she already know? Acting out the story in class should help you decide.

Get two volunteers to act as Mary and make a corner of the room as black-and-white as possible. You might give them a white tablecloth, a black book, a grey toy animal, or a plastic brain. You could dress them in white lab coats – whatever you have to hand. Ask the Marys to immerse themselves in the role of a futuristic colour scientist who knows *everything* physical a scientist can learn about the brain, the visual system, and colour: *everything*.

Now let the two Marys out in turn to do their best possible impersonations. ‘Mary-amazed’ acts completely surprised at what she sees, gasping at the delightful colours. ‘Mary-know-it-all’ explains why she is not surprised at all – how she understood everything in advance. Mary-know-it-all is the far harder role, so it may be best to choose someone who is familiar with the arguments. Sue once tried this at a Tucson conference only to discover afterwards that the volunteer Mary-know-it-all was Michael Beaton, inventor of RoboDennett – an unforgettable performance, especially as he was representing the argument he disagreed with!

Afterwards, everyone else can ask questions of the Marys, discuss their answers, and make up their own minds. Write down your own decision. You may find that it changes as you learn more about the nature of consciousness.

• SECTION ONE: THE PROBLEM

premises. We simply fail to follow the instructions – we fail to allow Mary to know *everything* physical there is to know about colour – because the instructions tell us to imagine something preposterously immense.

Dennett gives an alternative ending to the story. Mary's captors release her into the colourful world and, as a trick, present her with a blue banana. Mary is not fooled at all. 'Hey!' she says. 'You tried to trick me! Bananas are yellow, but this one is blue!' (1991, p. 399). She goes on to explain that because she knew *everything* about the physical causes and effects of colour vision, she already knew exactly what impressions yellow and blue objects would make on her nervous system, and exactly what thoughts this would induce in her. This is just what it means to have *all* the physical information about colour vision. When we readily assume that Mary will be surprised, it is because we have not actually followed the instructions – because it is really hard to imagine knowing absolutely everything physical about anything. And so, we have succumbed to 'Philosophers' Syndrome: mistaking a failure of imagination for an insight into necessity' (Dennett, 1991, p. 401).

To make it easier for us to imagine Mary having all the physical facts, Dennett (2005) invents 'a standard Mark 19 robot' with hardware equipped for colour vision, but with black-and-white video cameras installed instead of colour ones. While waiting for her cameras to be replaced, RoboMary learns all the physical information about the colour vision of Mark 19s:

She has all her hard-won knowledge of that system of color vision, but she can't use it to adjust her own hardware so that it matches that of her conspecifics. This doesn't faze her for a minute, however. Using a few terabytes of spare (undedicated) RAM, she builds a model of herself and *from the outside, just as she would if she were building a model of some other being's color vision*, she figures out just how she would react in every possible color situation.

(Dennett, 2005, p. 126)

And so, Dennett demonstrates that 'What RoboMary knows' leaves no space left over for her to be startled, delighted, or surprised.

British philosopher Michael Beaton retaliates with 'What RoboDennett still doesn't know'. Beaton argues that RoboMary cannot have a perfect model of herself, any more than RoboDennett can. And even if she could, she would be modelling the state of 'knowing what it is like' only as a state of the model, not of herself. Objective knowledge of oneself cannot necessarily be used as a simulation of oneself, and knowing all the facts about what one would say and how one would react is not the same as knowing all the facts about what it's like. Even if physicalism is true, Mary really could learn something new when she comes out (Beaton, 2005).

The imaginary Mary has led to many philosophical tangles like this (see e.g. Ludlow et al., 2004), and along the way her inventor even changed his mind about the argument against physicalism, suggesting that when we feel sure that Mary will learn something new we are under an illusion about our own experience (Jackson, 1998). Mary does not learn anything new, but merely finds herself in a different kind of representational state from those she was in before – that is, she now has the ability to recognise, imagine, and remember the state of seeing this colour (Jackson, 2003).

Some would conclude that this thought experiment, like many others, depends so precariously on linguistic hair-splitting – what counts as ‘knowing’ or ‘learning’, what do we mean by ‘physical information’, or indeed ‘all’? – that we can only ever get out of it what we put in. Others maintain that Mary can be helpful in making a tricky dichotomy easier to think about. If you believe that Mary will be surprised when she comes out, then you probably believe that consciousness, subjective experience, or qualia are something additional to knowledge of the physical world. If you think she will not be surprised, then you likely believe that knowing all the physical facts tells you everything there is to know – including *what it is like* to experience something.

THE PHILOSOPHER’S ZOMBIE

Imagine there is someone who looks like you, acts like you, speaks like you, and in every detectable way is identical with you, but is not conscious. An early form of this idea was the ‘zombie replica’, an exact physical replica of a man, ‘to which there applied *only* the physical descriptions and whatever they entailed’ (Kirk and Squires, 1974, p. 141), behaviour being one of the aspects entailed (Kirk, 2005). There have been many variations on this theme, but we will stick with the popular version proposed by Chalmers (1996), a creature that is indistinguishable from a conscious one. There is *nothing it is like* to be this creature. There is no view from within. No consciousness. No qualia. This – not something from *World War Z* or *The Walking Dead* – is the philosopher’s zombie.

This zombie has caused even more trouble than Mary. Most people agree that it is easy to imagine one, but is a zombie either logically or physically possible?

Chalmers thinks so: ‘the logical possibility of zombies seems [...] obvious to me. A zombie is just something physically identical to me, but which has no conscious experience – all is dark inside’. He goes on, ‘I can detect no internal incoherence; I have a clear picture of what I am conceiving when I conceive of a zombie’ (1996, pp. 96, 99). Chalmers’s zombie twin, living on zombie earth, is quite conceivable, he argues. He suggests we imagine a silicon version of Chalmers who is organised just like the real philosopher and behaves just like him but has silicon chips where the real one has neurons. Many people would expect such a creature to be unconscious (whether or not it would be). Then, he suggests, just replace the chips with neurons, and you have my zombie twin – totally indistinguishable from the real philosopher, but with nothing it is like to be him. This works, he argues, because there is nothing in either silicon or biochemistry that conceptually entails consciousness.

‘conscious states are qualitative states right down to the ground. Take away the qualia and there is nothing there.’

(Searle, 1998, p. 21)

‘A zombie is just something physically identical to me, but which has no conscious experience – all is dark inside’

(Chalmers, 1996, p. 96)



PROFILE 2.1

David Chalmers (b. 1966)



Born in Australia, David Chalmers originally intended to be a mathematician, but then he spent six months hitchhiking around Europe on his way to taking up a Rhodes

scholarship at Oxford, and spent most of his time thinking about consciousness. This led him to Douglas Hofstadter’s research group and a PhD in philosophy and cognitive science. He is responsible for the distinction between the ‘easy problems’ and the ‘hard problem’ of consciousness, and he is one of that rare breed: a self-proclaimed dualist. As well as helping to get a science of consciousness off the ground, his other interests include artificial intelligence and computation, philosophical issues about meaning and possibility, and the foundations of cognitive science. Working for many years in Tucson, Arizona, he founded the ‘Toward a Science of Consciousness’ conferences. He is now Distinguished Professor of Philosophy and director of the Centre for Consciousness at the Australian National University in Canberra, as well as Professor of Philosophy and co-director of the Center for Mind, Brain, and Consciousness at New York University.

'I have a clear picture of what I am conceiving when I conceive of a zombie'

(Chalmers, 1996, p. 99)

But if you think that consciousness has effects or functions, you will disagree with Chalmers. For example, if you believe we need to be conscious to think, speak, or make difficult decisions, then a creature without consciousness would be unable to do these things. This means it could not be indistinguishable from a conscious person, so zombies could not exist. Another way of saying this is that if zombies are possible, then consciousness must be superfluous, a kind of epiphenomenon that exists but does nothing. This is the idea of 'consciousness inessentialism'.

Imagine zombie earth: a planet just like ours, peopled by creatures who behave exactly like us, but who are all zombies. There is nothing it is like to live on zombie earth. In 'Conversations with Zombies', philosopher Todd Moody (1994) uses the following thought experiment designed to reject consciousness inessentialism. He imagines the whole zombie earth to be populated by people who use such terms as *think, imagine, dream, believe, or understand*, but who cannot understand any of these terms in the way we do because they have no conscious experience. For example, they might be able to talk about sleep and dreaming because they have learned to use the words appropriately, but they would not have experiences of dreaming as we do. At most they might wake up to a sort of coming-to-seem-to-remember which they learn to call dreaming.

On such an earth, Moody argues, the zombies might get by using our language, but zombie philosophers would be mightily puzzled by some of the things we conscious creatures worry about. For them the problem of other minds, or our worries about qualia and consciousness, would make no sense. They would never initiate such concepts as consciousness or dreams, so zombie philosophy would end up quite different from ours. From this he argues that although the zombies might be individually indistinguishable from conscious creatures, they would still show the mark of zombiehood at the level of culture. At this level, consciousness is not inessential – it makes a difference.

Moody's thought experiment inspired a flurry of objections and counter-arguments from philosophers, psychologists, and computer scientists (Sutherland, 1995). One of the main objections is that Moody has broken the rules of the thought experiment. It is worth reminding ourselves what exactly those rules are.

Chalmers's core definition concerns the physical: 'someone or something physically identical to me (or to any other conscious being), but lacking conscious experiences altogether' (1996, p. 94). But this entails behavioural identity too: 'my zombie twin is by definition physically identical to me over its history, so it certainly produces indistinguishable behavior' (1996, p. 120). This means that the people on zombie earth must be truly and wholly indistinguishable in all their actions, too. If their philosophy, or the terms they invented, were different, then they would be distinguishable from us and hence not count as zombies. If you really follow the rules, there is nothing left of the difference between a conscious human and a zombie.

Then again, maybe Moody's argument does exactly what a thought experiment is meant to do: helps us see something that wasn't already totally obvious. If you imagine a physically identical zombie and ask yourself what zombie culture would be like,



FIGURE 2.4 • Which is which? Can you tell? Can they?

and simply cannot imagine it being identical to our culture, perhaps that does tell you something.

Some philosophers think the whole debate is misguided. Patricia Churchland calls it 'a demonstration of the feebleness of thought-experiments' (1996, p. 404). Dennett thinks it is based on bogus feats of imagination. As they point out, being able to say that you can imagine something counts for nothing. If you know no science, you might say you could imagine water that was not made of H₂O, or a hot gas whose molecules were not moving fast. But this would tell us more about your ignorance than about the real world. Chalmers (2010) disagrees, arguing that we might conceive of a situation (say a twin-earth world) in which water is still H₂O but in which there is also a watery stuff that is not H₂O. The twin earth is metaphysically possible and is accessed by the act of conceiving. He distinguishes between different forms of conceivability and possibility and defends the legitimacy of using one as a guide to the other.

This debate goes right to the heart of how we perform thought experiments and why. But even those who are sceptical about stepping from conceivability to possibility or necessity, as Dennett is, continue to find the thought experiment a tempting tool. To help us think more clearly about zombies, Dennett introduces the concept of the zimbo. Imagine there is a simple zombie: some sort of creature (biological or artificial) that can walk about and behave in simple ways appropriate to its needs. Now imagine a more complex kind of zombie. In addition, this complex zombie also

monitors its own activities, including even its own internal activities, in an indefinite upward spiral of reflexivity. I will call such a reflective entity a *zimbo*. A zimbo is a zombie that, as a result of self-monitoring, has internal (but unconscious) higher-order informational states that are about its other, lower-order informational states.

(1991, p. 310)

Imagine a conversation with such a zimbo. For example, we might ask the zimbo about its mental images, or about its dreams or feelings or beliefs. Because it can monitor its own activities, it could answer such questions – indeed, it would do so in ways that would seem quite natural to us, and would suggest that it was conscious just like us. As Dennett concludes, 'the zimbo would (unconsciously) believe that it was in various mental states – precisely the mental states it is in position to report about should we ask it questions. *It* would think it was conscious, even if it wasn't!' (p. 311). This is how Dennett comes to make his famous claim that 'We're all zombies. Nobody is conscious – not in the systematically mysterious way that supports such doctrines as epiphenomenalism!' (p. 406). What he means is that we are complex self-monitoring zombies – zimboes – who can talk and think about mental images, dreams, and feelings; who can marvel at the beauty of a sunrise or the light rippling in the trees. But if we think that being conscious is something separable from all of this, we are mistaken. On this view there is no fundamental difference between phenomenal and access consciousness (Dennett, 1995a).

At its simplest, the zombie debate amounts to this. On the one hand, if you believe in the possibility of zombies, then you believe that consciousness is distinct from

'I take this argument to be a demonstration of the feebleness of thought-experiments'

(Churchland, 1996, p. 404)

'people can, without conceptual inconsistency, think the "impossible" thought that H₂O is not water'

(Papineau, 2003a, p. 361)

'The philosopher's debate on zombies is really just the qualia wars.'

(Sutherland 1995, p. 312)

the physical body and is an inessential optional extra to behaviour (this is epiphenomenalism or conscious inessentialism). We might do everything we do either with or without it and there would be no obvious difference. It is therefore

a mystery why we have consciousness at all. On the other hand, if you believe that zombies are not possible, you might be a dualist who believes we need a soul or non-physical mind as well as a body. But if you want to avoid dualism you must conclude that anything made like us, and behaving like us, would necessarily be conscious. The mystery in this case is not why we have consciousness at all, but why or how consciousness necessarily comes about in creatures like us. There are many different views in each of these camps (for a review, see Kirk, 2015), but this is the essential distinction.



THE PHILOSOPHER'S ZOMBIE

The most common form of the philosopher's zombie is defined by two statements.

- 1 The zombie is physically and behaviourally indistinguishable from a conscious human being.
- 2 There is nothing it is like to be a zombie. That is, the zombie is not conscious.

When thinking about zombies it is cheating if you allow your zombie to do things we would never do, or behave in ways we would not (then it would not fit statement 1). Equally, your zombie cannot have little bits of inner experiences or a stream of consciousness (then it would not fit statement 2). Most people agree that zombies are conceivable, but could they really exist?

- 1 If you say yes, then you believe that consciousness has no effects or consequences; it is an inessential extra and we could be just as we are and do everything we do without being conscious.
- 2 If you say no, you believe that we could not be as we are or do everything we do without consciousness; any creature like us would necessarily be conscious.

It is worth thinking very carefully about this and writing down your own answer now – Yes or No. You may change your mind as you learn more about consciousness, and you will encounter the zombie again.

Zombies appear in arguments about the hard problem (this chapter), the function and evolution of consciousness (Chapters 10 and 11), and artificial consciousness (Chapter 12).

CONCEPT 2.1

IS THERE A HARD PROBLEM?

We can now return to Chalmers's hard problem with more mental tools at our disposal. The distinction between the hard and the easy problems of consciousness relates directly to Nagel's question 'what is it like to be a bat?' and gets at the central issues of the two thought experiments just described: 'Why aren't we all zombies?' and 'What does Mary gain when she emerges from her black-and-white room?' The way people react to these thought experiments is intimately related to how they deal with the hard problem of consciousness.

At the risk of oversimplifying, we shall divide responses to the hard problem into five categories.

1 THE HARD PROBLEM IS INSOLUBLE

William James wrote long ago about believers in the soul and positivists who wish for a tinge of mystery. They can, he said, continue to believe 'that nature in her unfathomable designs has mixed us of clay and flame, of brain and mind, that the two things hang indubitably together and determine each other's being, but how or why, no mortal may ever know' (1890, i, p. 182).

More recently, the 'new mysterians' have argued that the problem of subjectivity is intractable or hopeless. Nagel, for example, argues that not only do we have no solution, we do not even have a conception of what a physical explanation of a mental

phenomenon would be. (This is reminiscent of one of the main objections to the thought experiment about Mary: how can we actually conceive of knowing all the physical facts about vision?) Colin McGinn (1999) describes the problem as a ‘yawning conceptual divide’ (p. 51), an irreducible duality in the way we come to learn about mind and brain. As he puts it:

You can look into your mind until you burst, and you will not discover neurons and synapses and all the rest; and you can stare at someone’s brain from dawn till dusk and you will not perceive the consciousness that is so apparent to the person whose brain you are so rudely eye-balling.

(1999, p. 47)

He claims that we are ‘cognitively closed’ with respect to this problem – much as a dog is cognitively closed with respect to reading the newspaper or listening to poetry. However hard the dog tried, it would not be able to master mathematics, because it does not have the right sort of intelligence. Similarly, our human kind of intelligence is wrongly designed for understanding consciousness. In McGinn’s view, we can still study the neural correlates of conscious states (what Chalmers would call one of the easy problems), but we cannot understand how brains give rise to consciousness in the first place.

‘our intelligence is wrongly designed for understanding consciousness.’

(McGinn, 1999, p. xi)

Psychologist Steven Pinker is equally defeatist. He thinks we can still get on with the job of understanding how the mind works, but our own awareness is ‘the ultimate tease [. . .] forever beyond our conceptual grasp’ (1997, p. 565).

‘the new mysterianism is a postmodern position designed to drive a railroad spike through the heart of scientism’

(Flanagan, 1992, p. 9)

Although the new mysterianism, unlike that of James’s day, is a naturalistic position rather than a supernaturalist one, it has also been described as a fundamentally postmodern challenge to the belief that science will eventually explain the whole of the natural world (Flanagan, 1992, p. 9). Thinkers in this category, who also include philosopher Jesse Prinz, all agree that there is a hard problem, and agree that we will never solve it.

2 TRY TO SOLVE IT

Some theorists believe that the problem is really hard but still soluble. Trying to solve the hard problem may, however, involve first restating it in different words. For example, Gray redefines the problem as ‘*How does the brain create qualia?*’ (2004, p. 301) or more specifically, ‘*how does the unconscious brain create and inspect the display medium of conscious perception?*’ (p. 123). These questions, particularly the second, limit in advance the range of possible solutions to the problem, as we will see when (in Chapter 8) we look at Gray’s theory as an attempt to answer them.

Others argue that a solution requires some fundamental new understanding of the universe – what Pat Churchland calls ‘a real humdinger of a solution’ (1996, p. 40). We have already met Libet’s conscious mental field, which he deemed necessary because ‘a knowledge of nerve cell structures and functions can never, in itself, explain or describe conscious subjective experience’ (2004, p. 184). And as we have seen, Chalmers’s own (1995a, 1996, 2007) attempt at a solution is a kind of dualism: a dual-aspect theory of information in which all information has two basic aspects, physical and experiential. So, whenever there is conscious experience, it is one aspect of an information state, and the other aspect lies in the

'The hard problem is a hard problem, but there is no reason to believe that it will remain permanently unsolved'

(Chalmers, 1995a, p. 218)

physical organisation of the brain. On this view, we will understand consciousness only when we have a new theory of information.

Others appeal to fundamental physics or to quantum theory. For example, the British mathematician Chris Clarke (1995) treats mind as inherently non-local, like some phenomena in quantum physics. In his view, mind is *the* key aspect of the universe and emerges prior to space and time: 'mind and the quantum operator algebras are the enjoyed and contemplated aspects [i.e. the subjective and objective aspects] of the same thing' (1995, p. 240). Chalmers's and Clarke's are both dual-aspect theories and are close to panpsychism.

The British mathematician Roger Penrose (1989) argues that consciousness depends on non-algorithmic processes – that is, processes which cannot be carried out by a digital computer, or computed using describable procedures (Chapter 12). With anaesthesiologist Stuart Hameroff, he has developed a theory that treats experience as a quality of space-time and relates it to quantum coherence in the microtubules of nerve cells (Hameroff and Penrose, 2014) (Chapter 5).

All these theories assume that the hard problem is soluble, but only with a fundamental rethink of the nature of the universe.

3 TACKLE THE EASY PROBLEMS

There are many theories of consciousness that attempt to answer questions about attention, learning, memory, or perception, but do not directly address the question of subjectivity. Chalmers (1995b) gives as an example Crick and Koch's theory of visual binding. This theory uses synchronised oscillations to explain how the different attributes of a perceived object become bound together to make a perceptual whole (Chapter 6). 'But why', asks Chalmers, 'should synchronized oscillations give rise to a visual experience, no matter how much integration is taking place?' (p. 64). Synchronised oscillations are offered as an 'extra ingredient' (1995a), but why should that particular ingredient account for consciousness? He concludes that Crick and Koch's is a theory of the easy problems. If you are convinced, as Chalmers is, that the hard problem is quite distinct from the easy problems, then many if not most theories of consciousness are like this, including theatre metaphors of processing capacity and attention (Chapters 5 and 7), evolutionary theories based on the selective advantages of introspection or the function of qualia (Chapter 11), and theories that deal with the neural correlates of consciousness (Chapter 4). In all these cases one might still ask: 'But what about subjectivity? How does this explain the actual phenomenology?'

Crick and Koch themselves say that the most difficult aspect of the problem of consciousness is the problem of qualia. From one perspective this sounds like a tautology: the most difficult thing about the problem of consciousness is the problem of consciousness. From their perspective, however, it makes perfect sense to split the problem up into harder and easier bits, and tackle the easier bits first. 'The history of the past three millennia has shown', they say, 'that it is fruitless to approach this problem head-on.' So instead of carrying on trying to explain how

the painfulness of pain or the redness of red arises from, or is identical to, the actions of the brain [...] we are attempting to find the neural

correlates of consciousness (NCC), in the hope that when we can explain the NCC in causal terms, this will make the problem of qualia clearer.

(2003, p. 119)

Hunting for the NCCs is currently one of the most popular ways of scientifically studying consciousness, and solving the hard problem is often presented as its ultimate aim. One recent argument is that two common methods – using brain scans to compare conscious and unconscious states, and to investigate specific ‘contents’ of consciousness – cannot meaningfully be used separately. Instead, future research should combine the two in order to measure ‘the relative contribution of the mechanistically distinguishable subcomponents of the brain involved in producing the astonishingly rich and often heartbreakingly [sic] beautiful phenomenal view of the world’ (Bachmann and Hudetz, 2014, p. 10). But this still does not explain *how* any method for finding more detailed correlations between neural activation and experience can be expected to bridge the gap. Invoking the hard problem may sometimes be a way of sprinkling a little of the glamour of the ‘beautiful phenomenal view’ on to fashionable neuroscientific research. In such cases, the researchers may think they are tackling the hard problem, but others might say they are tackling the easy problems as though they were the hard one.

‘it is the “easy” problem that is hard, while the hard problem just seems hard because it engages ill-defined intuitions’

(Dehaene, 2014, p. 262)

By contrast, French neuroscientist Stanislas Dehaene claims that we have it all the wrong way round:

My opinion is that Chalmers swapped the labels: it is the ‘easy’ problem that is hard, while the hard problem just seems hard because it engages ill-defined intuitions. Once our intuition is educated by cognitive neuroscience and computer simulations, Chalmers’s hard problem will evaporate [. . .] the science of consciousness will keep eating away at the hard problem until it vanishes.

(2014, p. 262)

‘there is no real distinction between hard and easy problems of consciousness, and the illusion that there is one is caused by the pseudo-profundity that often accompanies category mistakes’

(Pigliucci, 2013)

Others again would say it is mistaken to think that the hard and the easy problems are separable at all.

4 IDENTIFY MORE HARD PROBLEMS

Looking for the neural correlates of conscious experience raises many interesting questions about principles and methods, which we will consider in Chapter 4. One is tackled by the psychiatrist and neuroscientist Steven Miller, who argues that researchers working on NCCs often imply that finding them will help us identify the neural *constitution* of consciousness but fail to recognise that not every neural correlate of a ‘conscious state’ is necessarily constitutive of that state. That is, things might be going on in the brain that reliably accompany conscious experience but are not identical with it, and may even have nothing to do with causing it. This means that any given experience may be caused by more than one pattern of brain activity, and one pattern of brain activity may cause many different experiences. Understanding these relationships may conceivably be within science’s reach, but that does not necessarily mean these are ‘easy problems’. Miller asks, ‘might neural multiple-realizability problems be nevertheless equally hard problems to answer despite their being more easily conceived as scientific problems?’ (2007, p. 167).

What is it like to be . . . ?

Chapter Two

Miller goes on to split Chalmers's original hard problem into two: the *hard existence* problem (why and how do we have phenomenal consciousness at all?) and the *hard character* problem (why does particular brain activity feel like this and not like that?). He claims that multiple realizability might help us 'sharpen' the hard character problem, but not to solve it. He then considers related problems that may also be genuinely 'hard': the problem of direct intersubjective exchange (how to compare the redness or happiness experienced by two people), and the problems of ontogeny (e.g. when does consciousness arise in the development from zygote to embryo to foetus to baby) and phylogeny (when it arises in evolution).

In Chapter 5, we will meet some more variations on the hard problem, including computer scientist Scott Aaronson's 'Pretty-Hard Problem' (2014) of which physical systems are conscious and which are not. Chalmers splits this new problem up into four more, including 'PHP1', the problem of constructing a theory that matches our intuitions about which systems are conscious, and 'PHP4', the problem of constructing a theory that tells us which systems have which states of consciousness.

This leaves us with a lot of hard problems, and a lot of 'easy' territory that suddenly looks very slippery.

5 THERE IS NO HARD PROBLEM

Adopting a more gung-ho optimism, in 'There is no hard problem of consciousness' Kieron O'Hara and Tom Scutt (1996) give both methodological and philosophical reasons for ignoring the hard problem. First, we know how to address the easy problems and should start with them. Second, solutions to the easy problems will change our understanding of the hard problem, so trying to solve the hard problem now is premature. A solution to the hard problem would only be of use if we could recognise it as such, and for the moment the problem is not well enough understood: indeed, 'all discussion of [the hard problem] seems to preclude any sort of answer being given' (p. 291).

Confusions about consciousness are explored by the British philosopher David Papineau, who also warns us against trusting too blindly in our intuitions – in this case, those which tell us that the magic of consciousness arises from the soggy of grey matter. According to Papineau, we are seduced into thinking materialism is false because the concepts and terms we use to refer to brain states do not involve experiences in the way that words we use when we talk about mental states or feelings do. We have no problem with accepting that temperature and mean kinetic energy are just two ways of referring to the same thing, and we should do the same with pain and nociceptive-specific neuronal activity. The problem, he suggests, is that instead

We focus on the left-hand side, deploy our phenomenal concept of pain (that feeling), and therewith feel something akin to pain. Then we focus on the right-hand side, deploy our concept of nociceptive-specific neurons, and feel nothing (or at least nothing in the pain dimension – we may visually imagine axons and dendrites and so on). And so we conclude that the right hand side leaves out the feeling of pain itself, the unpleasant what-it's-likeness, and refers only to the distinct physical

correlates of pain. [. . .] There is no reason why we shouldn't be able to refer to this 'what-it's-likeness' using concepts which don't actually give us the feeling.

(2003b, p. 6)

In other words, we expect too much of the language we use to talk about the physical side of the consciousness equation, and this blinds us to the fact that it is an equation: that the physical activity *equals* the experience.

It may help to think in terms of two fallacies, Papineau suggests. The first, which we may be familiar with from Romantic poetry, is the 'pathetic fallacy', in which we attribute human feelings to nature – dramatic storm clouds reflect a stormy mood, for example. The mistake we make when thinking about consciousness is the opposite, the 'antipathetic fallacy', in which we fail to recognise that feelings exist in parts of nature, such as brains. If we could stop committing the antipathetic fallacy, then we would be able to accept the reality of materialism, and the hard problem would melt away. The only thing that stands in the way of solving the hard problem is an explanation of why materialism should *seem* false, even though it is true. So, this position is somewhere between reformulating the hard problem and denying that it exists.

Analytic philosopher Patricia Churchland goes further. The hard problem is 'ridiculous', she says (in Blackmore, 2005). It's a 'hornswoggle problem' (Churchland, 1996) – a grand hoax. First, we cannot, in advance, predict which problems will turn out to be easy and which hard. For example, biologists once argued that to understand the basis of heredity we would have to solve the protein-folding problem first. In fact, base-pairing in DNA provided the answer, and the protein-folding problem remains unsolved. So how do we know that explaining subjectivity is so much harder than the 'easy' problems? Also, she questions whether the 'hard' things – the qualia – are well enough defined to sustain the great division. For example, do eye movements have eye-movement qualia? Are there thought-qualia, or does thinking have the qualia of auditory imagery or talking to oneself? If things become so hazy so soon after we leave behind the usual cases of seeing the blue sky or feeling a brick land on our foot, perhaps the great gulf is narrower than it seems. Finally, the distinction depends on the false intuition that if perception, attention, and so on were understood, there would necessarily be something else left out – the something that we have and a zombie does not.

Dennett likens the argument to that of a vitalist who insists that even if all the 'easy problems' of reproduction, development, growth, and metabolism were solved, there would still be the 'really hard

'There is no reason why we shouldn't be able to refer to this "what-it's-likeness" using concepts which don't actually give us the feeling'

(Papineau, 2003b, p. 6)



PROFILE 2.2

**Patricia Smith Churchland
(b. 1943)**



Pat Churchland is best known for her books on neurophilosophy attempting to unify the mind–brain, and for

her outspoken views on the philosophy of mind. She advocates eliminative materialism, and her motto is 'To understand the mind, we must understand the brain'. She grew up on a poor but beautiful farm in British Columbia, where her parents were pioneers – of literal new territories. She is now Professor of Philosophy Emerita at the University of California, San Diego. She is married to the philosopher Paul Churchland, and they work closely together. She thinks the hard problem is a 'Hornswoggle problem' which will go the way of phlogiston or caloric fluid, that zombies demonstrate the feebleness of thought experiments, and that quantum coherence in microtubules is about as good a theory as pixie dust in the synapses.

'Chalmers's "Hard Problem" is a theorist's illusion'

(Dennett, 2005, p. 134)

problem: life itself' (1996a, p. 4). 'Chalmers's "Hard Problem" is a theorist's illusion [. . .] not a real problem to be solved with revolutionary new science' (2001a, p. 223; also 2005, pp. 134–135). When asked 'But what about the *actual* phenomenology?', Dennett replies: 'There is no such thing' (1991, p. 365). This is not because he denies that we are conscious, but because he thinks we misconstrue consciousness. It only *seems* as if there is actual phenomenology – what we need to explain is not the phenomenology itself but how it comes to seem this way.

These are versions of what, in consciousness studies, has come to be called 'illusionism' (see Chapter 3). This is 'the view that phenomenal consciousness, as usually conceived, is illusory' (Frankish, 2016b, p. 11). For the illusionist, what needs to be explained is not phenomenality or qualia or 'the experiences themselves' but our illusory ideas about experience. This means we can avoid the hard problem altogether and replace it with 'the illusion problem'.

All possible responses to the hard problem are still hotly disputed, and there are many more theories than we have mentioned here (for helpful reviews, see Seth [2007] and Seager [2016]). There is no doubt that the idea of subjectivity – what it's like to be – lies at the heart of the problem of consciousness. But beyond that there is plenty to doubt.



Churchland, P. S. (1996). The Hornswoggle problem. *Journal of Consciousness Studies*, 3, 402–408.

Dissects various bad reasons why we might put consciousness in a different class from all other problems.

Kirk, R. (2015). Zombies. In E. N. Zalta (Ed.), *The Stanford encyclopedia of philosophy* (Summer 2015 edition). <http://plato.stanford.edu/archives/sum2015/entries/zombies/>

Outlines the arguments for and against zombies and how zombies relate to conceivability and possibility, mental causation, and the function of consciousness.

Ludlow, P., Nagasawa, Y., and Stoljar, D. (Eds.) (2004). *There's something about Mary: Essays on phenomenal consciousness and Frank Jackson's knowledge argument*. Cambridge, MA: MIT Press.


Groups responses to Jackson (and his later replies) into categories: what, if anything, does Mary learn, and could she really know everything physical?

Nagel, T. (1974). What is it like to be a bat? *Philosophical Review*, 83, 435–450.

Approaches consciousness through the problem of other minds. Exploring the obstacles to physicalism, Nagel suggests an 'objective phenomenology' to help understand subjectivity.

Wright, E. (Ed.) (2008). *The case for qualia*. Cambridge, MA: MIT Press.

The nineteen chapters include philosophical and scientific defences and attacks. Students might be asked to read first section of the introduction (pp. 1–3) plus one additional chapter to present or discuss in class (chapter summaries are provided on pp. 22–42).



The grand illusion

CHAPTER

THREE

Most of the films we love watching are peppered with continuity errors great and small, but how many of us ever noticed a white van driving into a *Braveheart* battle scene, or bullet holes in a wall before shots were fired in *Pulp Fiction*? Should we doubt our perceptual grasp on the world?

The closer you are to something, the easier it is to feel that you understand it. Nothing, it seems, could be closer to me than my personal experience; after all, it is a large part of what makes me me. But in the first two chapters we have already begun to find that our intuitions about the nature of experience and its relation to our physical world and body may not always be completely reliable. And sometimes, perhaps if we have spent too long imagining Mary coming out of her black-and-white room over and over again, we lose hold of what those intuitions even are.

This can be a frightening moment: if I cannot base my exploration of consciousness on what I know about my own consciousness, what can I base it on? But it can also be liberating: OK, now it makes sense for me to go right back to the beginning and work my way back to the question I was struggling with, one careful step at a time. One crucial part of doing this is being willing to accept the possibility that you are mistaken about some aspect of your experience. In this sense, we must be prepared to ask whether some of the ways things seem to us might be illusions.

The word 'illusion' is sometimes taken to mean something that doesn't exist: 'His arrogance was just an illusion.' But more precisely, an illusion is something that it is not what it appears to be: what looked like arrogance was profound shyness.

But the distinction between something being non-existent and being other than it seems is tricky, because once you say something is not what it seems to be, you may decide you need a new word for it, and so you do end up replacing the old with the new – that is, declaring that the old thing does not exist. You will notice these ambiguities popping up in many cases where consciousness, or free will, or reality, are called ‘illusions’.


‘The term illusion instantly aligns people’s thoughts in the wrong direction’

(Graziano, 2016, p. 112)

Some of the most familiar things that might spring to mind when we think of illusions are visual illusions. In Figure 3.1, for example, the lines and shapes really do exist, but the pyramid you see is an illusion: there is something there, but it’s not a pyramid. Applying the same idea to consciousness, we might say, as illusionists do, that our experiences exist but consciousness, in the sense that many people imagine it, does not. It is no coincidence that the science of visual perception is one of the areas in which the idea of illusion has become most important. Vision has been more thoroughly studied than any other sensory modality, and it is also the sense that to many people feels more essential to consciousness than any other: when I think about what it’s like to be me, the visual experience of looking out and seeing the world may well be the first thing that comes to mind. Where there is competition between the senses, vision usually trumps the others, although for people with little or no vision, hearing is often the most dominant sense from which they construct their understanding of the world. Finally, vision also has the special status of being more closely associated with knowledge than any other. Our ordinary language is full of metaphors that make seeing equivalent to knowing: ‘I see what you mean’, ‘her argument was crystal clear’, ‘we have looked carefully into the evidence’. These associations may make it uncomfortable to admit that our visual sense might be misled or misleading in some way. They also mean that in the case of vision it is all the more important to consider this possibility, in case our strong intuitions turn out to be false.

Many familiar illusions, visual and otherwise, may be handy shortcuts that for many everyday purposes work perfectly well – like assuming the world is flat when we’re driving, or that the sun is actually going down when we’re admiring a sunset. But trying to see things – even things as complex as vision itself – as they are, rather than as they seem, is crucial if we want our theories of consciousness to be theories of what consciousness is actually like rather than what we first leap to assuming it is like.

So, in this chapter we will take vision as a central example of how our conscious experience may be subject to illusions, in the hope that the idea of illusion may be a useful guide through one part of the labyrinth.



a pitcher and the wash basin – or a corner of the room with the table and the coat rack – appeared so not-real to me, despite their indescribable ordinariness, so completely not real, sort of ghostly, and at the same time provisional, waiting, temporarily taking

FIGURE 3.1 • Can you see a pyramid? If so, you are experiencing an illusion. An illusion is not something that does not exist, but something that is not what it appears to be.

up, as it were, the space of the real pitcher, the real wash basin filled with water

(Hugo von Hoffmansthal, 'Letters of the returning one' [Briefe eines Zurückgekehrten] IV, 26 May 1901, our translation)

First, let us go back to the beginning. What is it like to see? In particular, what is it like to have a conscious visual experience, such as consciously seeing a yellow daffodil on a green lawn? You see it; you can reach out to it; you delight in the rich visual experience of the petals' bright translucent yellow against the blades of green. Choose something in your field of view, and really look at it: be consciously aware of the curve of the tea cup or the pattern of the carpet. What does it mean for you to be seeing this, right now? What is it really like?

Seeing comes so naturally that these questions may seem silly, but they are not. Indeed, the difficulty of answering them has led some to conclude that visual experience is all a grand illusion. The term 'grand illusion' (Noë et al., 2000; Noë, 2002) emerged from research on change blindness and inattention blindness (discussed later in this chapter) to convey the idea that our visual experience may not be quite how it first seems. What sort of illusion do they mean?

Simple visual illusions, such as the effects of illusory contours, brightness and colour constancy, or the Müller-Lyer and café-wall illusions, are tricks that mislead you about what is out there in the world; they create confusion between appearance and reality. The interesting possibility for students of consciousness is not that we are sometimes wrong about what we are seeing, which we clearly can be, but that we may be wrong about the nature of seeing itself.

'We [...] are the victims of an illusion – not a perceptual illusion about the world but rather an illusion about the nature of our visual experience'

(Noë et al., 2000, p. 100)

The starting point, then, is how vision *seems*. How does it seem to you? It is important, before we go any further, to answer this question for yourself. This is partly because sometimes people propose novel solutions to difficult problems only to find that others say, 'Oh, I knew that all along'; and partly because some of the debates over the grand illusion concern the difficulty of knowing how people's experience really seems to them. And we cannot decide whether we need to talk about illusions unless we first know how it seems. So – how does it seem to you?

Close your eyes, reopen them, and look around. It probably seems as though you see the world like a richly detailed and ever-changing picture; perhaps as you turn your head to see what's on either side of you it seems more like a moving picture, a continuous 'stream of vision'.

Now, before going any further, it may also be useful to describe to yourself how you think vision actually works; try to come up with a basic theory about what is going on. Perhaps you arrived at something like this:

When we look around the world, unconscious processes in the brain build up a more and more detailed representation of what is out there. Each glance provides a bit more information to add to the picture. This rich mental representation is what we consciously see at any time. As long as we are looking around there is a continuous stream of such pictures. This is our conscious visual experience.

There are at least three threads of theory here. The first is the idea that there is a rich array of conscious visual impressions to be explained. The second strand is that at any time there are definite contents of which we are aware, while everything else remains 'outside our conscious awareness'. This is what Dan Dennett (1991) rejects when he claims that there is no show in the Cartesian theatre, no time and place where things come on to the stage and thus become conscious (we will return to this concept in more detail in Chapter 5). The third strand is the idea that seeing means having internal mental pictures; i.e. that the world is *represented* in our heads.

All these ideas are combined in concepts like James's stream of visual consciousness (1890, i, p. 245), the 'movie-in-the-brain' (Damasio, 1999), or 'the vivid picture of the world we see in front of our eyes' (Crick, 1994, p. 159). The emphasis placed on the dynamic flow of experience varies in these metaphors along a spectrum from stream to movie to picture. In all of them, however, the richness of the experience is unquestioned, and the in/out distinction and underlying representation tend to be too. The standard scientific model of vision seems to be built on the same assumptions as the intuitive account – but maybe both need questioning.

The idea that in seeing (and imagining) we represent the world in our minds goes back at least as far as the ancient Greeks, who thought about vision in terms of the world being reflected in the pupil of the eye (and also thought about imagination



FIGURE 3.2 • The Café Wall illusion, first described by Richard Gregory after seeing tiles on the wall of a café in St Michael's Hill in Bristol. When the tiles are dark and white and the mortar is thick enough and mid-tone the horizontal lines do not look parallel. No attempts to convince yourself that they are gets rid of the illusion.



PRACTICE 3.1

HOW MUCH AM I SEEING NOW?

As many times as you can, every day, ask yourself '**How much am I seeing now?**'

Whether you are looking at a busy street or a beautiful garden, a page of text or the back of your own hand, ask *How much am I seeing?* You may at first get the impression that you can see everything at once; that there is an entire, detailed scene in your awareness. Now look again, harder. What are you actually seeing right *now*?

If you do this a few hundred times you may be better able to assess the various theories covered in this chapter. Eventually you may notice some profound changes. Can you describe what has happened?




FIGURE 3.3 • Descartes believed that pictures were transmitted through the eyes to the pineal gland where they entered the mind. His theory has generally been rejected but the idea of pictures in the head remains popular.

'It is important to avoid the temptation of thinking that eyes produce pictures in the brain'

(Gregory, 1966/1997, p. 5)

as a kind of picture-viewing); this led naturally to conceiving of pictures inside the eye and the head. Leonardo da Vinci compared the eye to a camera obscura – a dark chamber, popular at the time, into which an image of the world is projected. Then, in the early seventeenth century, Kepler explained the optics of the eye but said he would leave to others the job of explaining how the image 'is made to appear before the soul' (Lindberg, 1976). This is what Descartes tried to do. He studied actual images by scraping off the back of an ox's eye so that he could see them form on the retina and then showed, in his famous sketches, how he thought these images are transmitted to the non-material mind.

The details of Descartes' scheme were overthrown, but the idea of pictures in the head remained, and it was updated by cognitive psychologists in the twentieth century, who talked of internal screens or models as well as pictures. Dennett calls the idea of pictures in the head 'an almost irresistible model of the "end product" of visual perception', but also a 'ubiquitous malady of the imagination' (1991, p. 52). Alva Noë and the French psychologist Kevin O'Regan are similarly confident that 'The supposed fact that things appear pictorial to us in no way requires there to be pictures in the head' (2001,

p. 947). Indeed, they even challenge the idea that it really *seems* to us as though we are looking at a picture: 'it is just bad phenomenology to assert that we take ourselves to have a 3D-model or picture in the head when we see' (p. 962).

One of the tricky questions raised by picture-in-the-head theories is what the information in the picture is for: are there some structures in the brain which make up the picture and others which read off the information contained in it? This risks requiring a whole mind-within-a-mind, often referred to as a 'homunculus': a little person inside your head. And arguably the little person in your head looking at your pictures also needs one in its head looking at its internal pictures, so we have only pushed the required explanation back a level. So, if you do think that seeing must involve having a conscious stream of pictures like a high-definition movie playing in the head, you will not be alone, but could you be wrong?

Only from time to time the pupil's shutter

Will draw apart: an image enters then,

To travel through the tautened body's utter

Stillness – and in the heart to end.

(Rainer Maria Rilke, 'The Panther' [Der Panther], 1902, our translation)

To sum up, there are three assumptions made about vision in much of the scientific tradition of its study, and all three may also figure in our intuitions about vision: 1) visual experience is richly detailed, 2) there are things that are in and out of our visual experience, and 3) vision operates by representing the world in the mind or the brain. Perhaps these assumptions seem unremarkable. However, they can land

us in difficulty once we start trying to apply them to understanding how the brain contributes to visual, or any other, experience. If we assume all three, we end up having to explain how all the neural processing in all the parallel pathways in the human visual system results in that rich, definite, representation-based conscious experience. We also have to work out what it is that distinguishes all that mass of ‘unconscious’ processing from the final ‘conscious’ representation. What creates the ‘magic difference’ between some representations being conscious and others not?

One way to approach this question – which amounts to a version of the hard problem – is to stick with the idea of a stream of conscious visual representation and look for its neural correlates (Chapter 4). The basic principle is simple. If you believe that some visual representations in the brain are conscious and others are not, then you should be able to take examples of each and study them in detail until you discover the difference. In this light Francis Crick asks, ‘What is the “neural correlate” of visual awareness? Where are these “awareness neurons” – are they in a few places or all over the brain – and do they behave in any special way?’ (1994, p. 204). He goes on to consider synchronised behaviour in widely separated neurons (Chapter 6), but adds that ‘so far we can locate no single region in which the neural activity corresponds exactly to the vivid picture of the world we see in front of our eyes’ (p. 159).

Research on the neural correlates of vision has since progressed far enough that, using fMRI to build up a library of correspondences between what someone is viewing and their brain activity, scientists can infer backwards from new patterns of activity to the stimuli being perceived (Nishimoto et al., 2011; Poldrack, 2011). A similar technique has been used to match up brain activation with people’s reports of what they were dreaming about (Horikawa et al., 2013). But however good we get at finding these correspondences, remember they are finding the NCs of particular visual experiences, not the NCs of consciousness itself.

It is easy to imagine that this clever method is reading off the brain’s own internal pictures, but in fact it relies on complex patterns widely distributed across the cortex. Perhaps we should challenge the natural trio of ideas that conscious vision is as detailed as a hyper-realist painting or an HD 3D film, with things categorically in or out of the frame, and all dependent on picture-like representations.

FILLING IN THE GAPS

The perceptive William James noticed something very odd, although it is obvious once someone points it out: when we look around we do not, and cannot, take in everything at once, and yet we are unaware of any gaps. Imagine you have been sitting in your friend’s living room for an hour and suddenly notice that there is a vase of flowers on the table. What was there before? More wallpaper? A flower-shaped gap?

It is true that we may sometimes be tempted to exclaim, when once a lot of hitherto unnoticed details of the object lie before us, ‘How could we ever have been ignorant of these things and yet have felt the object, or drawn the conclusion, as if it were a *continuum*, a *plenum*? There would have been *gaps* – but we felt no gaps. . . .’

(1890, i, p. 488)

'There would have been gaps – but we felt no gaps'

(James, 1890, i, p. 488)

Why don't we notice the gaps? One answer might be that, in some sense, the brain fills in the missing bits: it pastes in more stripy wallpaper behind the vase, say. But if the brain already knows what needs to be filled in, who does it do the filling-in for, and why? Another answer is that there is no need to fill anything in because the gaps are just a lack of information. And an absence of information is not the same as information about an absence.

Consider something that happens in vision all the time: we infer the presence of whole objects from their visible parts. A car parked behind a tree looks like a whole car, not two halves separated by a tree trunk; a cat sleeping behind a chair leg looks like a whole cat, not two odd-shaped furry lumps. This ability to see objects as whole is obviously adaptive, but what is going on? We don't literally 'see' the hidden parts of the car, yet the car seems whole. This is sometimes referred to as amodal perception or conceptual filling-in. The car is conceptually completed but not visually filled in.

A more controversial kind of filling-in arguably happens in the blind spot. Where the optic nerve leaves the back of the eye there are no photoreceptors, creating a blind spot on the retina that subtends about 6 degrees of visual angle, roughly 15 degrees away from the fovea. Most people are unaware of their own blind spots until shown a demonstration such as that in Figure 3.4. Partly this is because we have two eyes, and the two blind spots cover different parts of the visual world, though even with one eye the blind spot is normally undetectable. But experiments can easily reveal it.




FIGURE 3.4 • Shut your right eye and look steadily at the small black spot with your left eye. Starting at about half a metre away, slowly bring the book towards you until the striped area disappears. It is then in your blind spot. You may need to try a few times to find it (remember to keep your eye on the black spot). What can you see there? Is the space filled in? Do the black lines join up? (Based on Ramachandran and Blakeslee, 1998, p. 95.)

A small object can be made to disappear from sight by lining it up precisely on the blind spot. What is seen where the object should have been? Not a blank space or a gaping black hole, but a continuation of the background. If the background is boring grey, then boring grey fills the space where the object should have been. If the background is black and blue stripes, the stripes seem to cover the whole area. The obvious conclusion is that the brain has somehow filled in the gap with grey or pink, or stripes or checks (or more people in the crowd, or more pebbles on the beach). But is this the right conclusion?

The American neuroscientist Christof Koch thinks so. 'Unlike electronic imaging systems, the brain does not simply neglect the blind spot; it paints in properties at this location using one or more *active processes* such as completion, interpolation, and filling-in' (2004, p. 54).

Dennett thinks not. This kind of thinking 'is a dead giveaway of vestigial Cartesian materialism' (1991, p. 344) – i.e. pretending to be a materialist but falling back into dualism. He challenges it with a thought experiment: imagine walking into a room papered all over with identical portraits of Marilyn Monroe. You can see at a glance (or a few glances) that the portraits are all the same. If one had a moustache, or a silly hat, you would notice straight away. It seems natural to conclude that you have seen the room in all its detail and now have a rich representation of it in your head.

This cannot be right. As Dennett points out, in order to identify one of the portraits you would have to look straight at it so that the image fell on the fovea.

While you did that all the other portraits would just be face-shaped blobs. Now you turn to the next, foveate that one, and again conclude that it is Marilyn and that it looks just the same as the first one. Now another . . . You can make at most four or five saccades a second. So, you cannot possibly check every single one in the brief time it takes you to conclude – ‘all Marylins’. You never see just one clear portrait and a lot of blurry blobs, you see the whole detailed lot of them. How can this be?

Could the brain be taking one of its high-resolution foveal views of the portrait and reproducing it lots of times, as if by photocopying, over an inner mental wall? Of course not, says Dennett. Having identified one portrait, and using texture-detection mechanisms to identify that all the blobs are of a similar size and shape, and finding nothing to suggest that the other blobs are not also Marylins, the brain jumps to the reasonable conclusion that the rest are Marylins too, and labels the whole region ‘more Marylins’. This is more like paint-by-numbers than filling in pixel-by-pixel. The reason you would notice a moustache or a silly hat is that you have dedicated pop-out mechanisms to detect such anomalies. If none of these are activated, the conclusion ‘all the same’ stands.

Of course, it does not seem that way to you. You are convinced that you are seeing lots of identical Marylins (or Dans, in our picture), and in a sense, you are. There are lots of portraits out there in the world, and that is what you are seeing. Yet it does not follow that there are lots of identical faces represented in your brain. Your brain just represents *that* there are lots: ‘no matter how vivid your impression is that you see all that detail, the detail is in the world, not in your head’ (Dennett, 1991, p. 355).

This does not apply only to the multiple Marylins room, or our multiple Dans picture. When you walk along the street, you cannot possibly look at all the detail around you, yet you see no gaps in the places where you haven’t looked. Does the brain fill in the spaces with plausible guesses about cars and trees and shop windows and children running to school? Does it need to?

There is a range of ideas about filling-in (Pessoa, Thompson and Noë, 1998; Komatsu, 2006). One, known as isomorphic filling-in, is that the brain actually fills in the details as though to complete a picture in the brain (or ‘in consciousness’). According to Koch, ‘the brain does not simply neglect the blind spot: it paints in properties at this location’ (2004, p. 54). Another, known as symbolic filling-in, is that the process is more conceptual than picture-like and occurs higher up the visual system. The most sceptical view is that there is no need to fill in anything at all.

Neuropsychologist V.S. Ramachandran reports both formal and informal experiments on a range of cases (Ramachandran and Blakeslee, 1998). With normal observers, if two vertical lines are shown, one above and one below the blind spot, the observer sees one continuous line. The lines can be offset slightly and still seem to form a single straight line, but if the same is done with horizontal lines they do not line up. Missing corners are not completed, but if the blind spot is positioned

‘We depicted consciousness as a place peopled with small imitations and these imitations were the images’

(Sartre, 1940/2004, p. 5)

- **Do not turn the page yet.** On the next page you will see an illustration. Try to look at it for just three seconds. You might like to practise counting at the right speed first, or get a friend to time you. Then turn the page, look at the picture while you count to three, and then turn back.

What did you see? Try to describe the picture in words before you look again.





FIGURE 3.5 • Perhaps you clearly saw lots of identical portraits of Dan Dennett, and just one with horns and a scar. But in three seconds you could not have looked directly at each one. Did you fill in the rest? Do you need to?



PROFILE 3.1

**Vilayanur S. Ramachandran
(b. 1951)**



Usually known as Rama, V. S. Ramachandran is a flamboyant neuroscientist and lecturer. Born in Tamil Nadu, he trained as a doctor in India, did a PhD at Trinity College, Cambridge, and then worked on visual perception and neurology. He is Director of the Center for Brain and Cognition, and Professor of Psychology and

over the centre of a radiating pattern – like a bicycle wheel with the centre left out – the pattern is completed and the lines are seen to converge to a point (Activity 3.1).

In one demonstration, Ramachandran uses a group of yellow doughnut shapes, with the central hole in one of them coinciding with the blind spot. A complete yellow circle appears and pops out from the surrounding doughnuts. From this he concludes that filling-in cannot be just a question of ignoring the gaps, because in that case the circle would not pop out. (A similar logic applies to experiments showing that synaesthesia involves visual rather than imaginative experience, which we come to in Chapter 6.) This finding shows, he says, that ‘your brain “filled-in” your blind spot with yellow qualia’ (Ramachandran and Blakeslee, 1998, p. 237). But what exactly are yellow qualia? Are they the same as Koch’s ‘properties’ that the brain ‘paints in’? Are they a form of Dennett’s fanciful ‘figment’, a non-existent kind of pigment used to paint in the blank space ‘in here’ (1991, e.g. pp. 346, 353)? If not, what is going on?

Other experiments used special participants, such as Josh, whose right primary visual cortex was penetrated by a steel rod in an industrial accident. He has a large permanent scotoma (or blind area) in his left visual field. Like other people with similar brain damage, he manages perfectly well for most purposes and, although well aware that he has a large blind spot, does not see a black hole or a space with nothing in it. 'When I look at you,' he said to Ramachandran, 'I don't see anything missing. No pieces are left out' (Ramachandran and Blakeslee, 1998, p. 98).

Ramachandran presented him with vertical lines above and below the large scotoma. At first Josh reported seeing a gap between the lines, but then the gap began to close, and he saw the lines growing slowly together until they met in the middle. Offset lines took about five seconds to line up and grow together. In other experiments, a column of large Xs was never completed across the scotoma, but a column of tiny Xs was. Ramachandran speculated that two different levels of visual processing were involved: the former activating temporal lobe areas concerned with object recognition, the latter treating the Xs as a texture and therefore completing them. Oddly enough, when a row of numbers was used, Josh reported that he could see numbers in the gap but could not read them, an odd effect that sometimes happens in dreams. Finally, when presented with a display of twinkling black dots on a red background, Josh reported that first the red colour bled into his scotoma, then the dots appeared, and last of all the dots began to twinkle. These results suggest not only a real effect, but one that takes a measurable time to occur, and can treat things like colour, texture, and movement separately.

The same was found with 'artificial scotomas' in unimpaired participants. Ramachandran and Gregory (1991) asked them to fixate the centre of a display of flickering 'snow' on a screen. Offset by 6 degrees was a small grey square with no snow. At first the square was visible to people, but after about five seconds it became filled with snow like the rest of the screen. When the whole screen was then made grey, a square of snow was seen and persisted for two to three seconds.

Experiments with monkeys showed increasing activity in area V3 corresponding to this effect (De Weerd et al., 1995). More recent investigation of how the blind spot is represented in monkeys' V1 found that the area dedicated to the blind spot is organised much like other parts of V1. This suggests that V1 contains a continuous functional topographic map rather than following the distribution of photoreceptors. So although the V1 maps are usually described as 'retinotopic' (representing the layout of the retinal image), they would better be described as 'visuotopic' (Azzi et al., 2015). In other words, what we see is not

Neuroscience at the University of California, San Diego. Ramachandran is best known for research on phantom limbs, rare neurological diseases, synaesthesia, and the 'bedroom intruder' in sleep paralysis. His original thinking has sometimes led to criticism that his speculations reach far beyond the evidence. He loves Indian painting and sculpture, and he thinks that the blind spot is filled in with qualia and that subjectivity resides mainly in the temporal lobes and cingulate gyrus.




FIGURE 3.6 • What happens this time – is there a gap or is it filled in with pebbles? If it is, are they large or small pebbles, or a random mix, or what?



ACTIVITY 3.1

Filling-in

With some simple experiments you can experience filling-in and explore its limits. In Figures 3.4 and 3.6, shut or cover your right eye and fixate the small black dot with your left eye. Hold the picture at arm's length and then move it gradually towards you until the larger circle disappears. Do you see a gap or a continuation of the background? Is the black line completed across the gap? What happens to the pebbles?

You can also try the effect with real people. It is said that King Charles II, who was a great promoter of science, used to 'decapitate' his courtiers this way. To do this in class, ask someone to stand in front while everyone else aims their blind spot at the victim's head. If you have trouble doing this, try the following. Hold up Figure 3.4 so that the circle disappears. Now, keeping the book at the same distance away from you, line up its top edge below the person's chin with the circle directly below. Now fixate whatever you can see above the black dot and remove the book. Your blind spot should now be on the person's head. Does the whole head get filled in? If not, why not? Does it matter how well you know the person?

You can explore what can and cannot be filled in by using your own pictures. Cut out a small fixation spot and a larger circle, or find suitably sized stickers, and stick them on. Alternatively fix them to a computer screen and experiment with moving displays. If you are doing several experiments it is worth putting a patch over your eye. With a stop-watch, you can time how long filling-in takes for different displays.

Can you deliberately prevent filling-in? Can you speed it up by making an effort? Does what you see in the gap ever surprise you? Can you explain the difference between those things that do and do not get filled in?

just what is on the retina, and this distinction is present early in the visual system. In humans, fMRI has been used to investigate the filling-in of contours, and several studies suggest that activity at the blind-spot region in V1 is closely linked to changes in perception, for example when what is seen alternates between two options in binocular rivalry (Meng and Tong, 2004). However, colour and brightness filling-in seem to be rather different from other kinds, with at least two studies finding no evidence that early visual areas (V1 and V2) contain map-like representations of brightness and colour that could be filled in (Perna et al., 2005; Cornelissen et al., 2006).

Filling-in also happens with afterimages, which are unique in retaining the same position on the retina when the eye moves. It seems that filling-in works differently for the original image and the afterimage, in the way the colours spread into each other (Hamburger, Geremek and Spillmann, 2012). This adds another dimension to evidence about whether the retina or the cortex is responsible for filling-in, suggesting that afterimages are treated by the brain as genuine stimuli that can create their own perceptual effects, rather than just obeying the rules of the retinal stimulation that first elicited them.

In other experiments, a background of English, Latin, or nonsense text was used. The square was filled in all right but, like Josh with his numbers, the participants could not read the text produced. This also throws doubt on the idea that filling-in is literally a process of completing a picture dot by dot. For how would one create visible letters and numbers that could not be read? But what then is it?

Contrary to the extreme sceptical view, these results clearly indicate that there is a real effect to be explained. The brain does not just ignore a lack of information, but responds in various ways at varying speeds. However, we cannot make sense of the findings by assuming that somewhere inside the brain there is a picture-like representation of the current object of perception, which must be filled in all over or gaps will be noticed.

Some kind of dynamic spreading of activation clearly does create illusory contours and the like. But this does not mean that this process is used to fill in an internal metric picture of the world. We might instead think of the blind spot, like other constants of our perceptual apparatus (such as how retinal resolution and colour

sensitivity drop off drastically towards the periphery) as being *used* in order to see. For example, 'if retinal sensation were not to change dramatically when an object falls into the blind spot, then the brain would have to conclude that the object was not being seen, but was being hallucinated' (O'Regan and Noë, 2001, p. 951). Like the curvature of the lens and the different functions of the rods and the cones, the blind spot is just one of the sensorimotor contingencies that shape our perceptual

experience. This predictable gap is not a problem for perception; it is an integral part of how it has evolved to work.

So how are gaps dealt with in perception more generally? One idea is to investigate where the mechanisms responsible for filling-in start to converge with those responsible for ordinary perception, to help identify ‘the critical stage at which subjective visual experience emerges’ (Komatsu, 2006). The argument is that because the two have different starting points – one beginning with photoreceptors responding to external stimuli, the other not – there must be a specific point at which the two converge and perceptual consciousness emerges. But this again commits us to some version of the idea that we can only experience what is represented somewhere ‘in consciousness’.

CHANGE BLINDNESS

Look at the picture in Figure 3.8 for a few moments. As you explore it you probably make many saccades and blinks, but you hardly notice these interruptions. It feels as though you look over the picture, take it in, and now have a good idea of what is there. Now ask yourself this question. If the tray under the teapot disappeared while you blinked, would you notice? Most people are sure they would.

Research showing they are wrong began with the advent of eye trackers, which made it possible to detect a person’s eye movements and make a change to a display during a saccade. In experiments beginning in the 1980s (Grimes, 1996), participants were asked to read text on a computer screen and then, during their saccades, parts of the surrounding text were changed. An observer would see the text rapidly changing, but the participants themselves noticed nothing amiss. Later experiments used complex pictures, with an obvious feature being changed during saccades. The changes were so large and obvious that under normal circumstances they could hardly be missed, but when made during saccades they went unnoticed.

This may seem very strange, but the effect is easily explained by the links between eyes and brain. Under normal circumstances, motion detectors quickly pick up transients and direct attention to that location. In set-ups like these, however, this mechanism is disabled. A saccade causes a massive blur of activity that swamps out these mechanisms, leaving only memory to detect changes. The implication is that trans-saccadic memory is extremely poor. With every saccade, most of what we see must be thrown away.

This research complements earlier work on trans-saccadic memory and visual integration across blinks and saccades (for a review, see Irwin, 1991). For a long time, it was taken for granted that the visual system must somehow integrate its successive pictures into one big representation that would remain stable across body movements, head movements, eye movements, and blinks. This would, of




FIGURE 3.7 • A field of yellow doughnuts. Shut your right eye and look at the small white dot near the middle of the illustration with your left eye. When the page is about six to nine inches from your face, one of the doughnuts will fall exactly around your left eye’s blind spot. Since the black hole in the centre of the doughnut is slightly smaller than your blind spot, it should disappear and the blind spot then is ‘filled in’ with yellow qualia from the ring so that you see a yellow disc rather than a ring. Notice that the disc ‘pops out’ conspicuously against the background of rings.

Paradoxically, you have made a target more conspicuous by virtue of your blind spot (Ramachandran and Blakeslee, 1998, p. 236).

‘Is the richness of our visual world an illusion?’

(Blackmore et al., 1995, p. 1075)



FIGURE 3.8 • When these two pictures are alternated with brief flashes of grey in between, or moved slightly when they are swapped, people rarely notice the change. This is one way to demonstrate change blindness.

course, be a massive computational task, and although it was not clear how it could be achieved, most researchers had assumed that somehow it must be – otherwise how could we have such a stable and detailed view of the world in consciousness? Change blindness challenged that assumption. Perhaps we do *not* have a stable and detailed view of the world at all, in which case massive integration of successive views is not necessary.

In fact, expensive eye trackers are not needed to induce change blindness. In the first experiment to use pictures, the effect was obtained by moving the image slightly (Blackmore et al., 1995). This forces the eyes to move, resembling what happens in a natural saccade. Subsequently, other less direct methods were developed, such as using image flicker, cuts in movies or during blinks, or a brief blank flash between pictures (Simons, 2000), and all seem to have similar effects (Domhoffer, Unema, and Velichkovsky, 2002). Motion detectors are also defeated by changes that are too slow to produce transients, and this provides another method for eliciting change blindness (Simons, Franconeri, and Reimer, 2000).

That the findings are genuinely surprising was confirmed by experiments asking people to predict whether they or others

would notice the changes under various conditions. Typically a large metacognitive error or ‘change blindness blindness’ was found – that is, people grossly overestimated their ability to detect change (Levin, 2002). It is this ‘discrepancy between what we see and what we think we see’ that justifies using the term ‘illusion’: ‘our awareness of our visual surroundings is far more sparse than most people intuitively believe’ (Simons and Ambinder, 2005, pp. 48, 44).

One of the simplest methods for demonstrating change blindness is the flicker method, developed by psychologists Rensink, O’Regan, and Clark (1997). They showed an original image alternating with a modified image (each shown for 240 msecs), with blank grey screens (for 80 msecs) in between, and counted the number of presentations before the participant noticed the change. With blanks in between, it took many alternations to detect the change; without the blanks it took only one or two.

They used this same method to investigate the effects of attention. When changes were made in areas of greater interest, an average of seven alternations was needed for participants to notice the change, whereas changes in areas of lesser

interest took an average of seventeen alternations, with some participants taking up to eight alternations to notice a change that was obvious once seen. This suggests that for unattended parts of the image, people have to do a slow serial search to find the change.

But even highly salient features can be subject to change blindness. For example, very gradual changes in facial expression can go unnoticed while still affecting people's subsequent behaviour (Laloyaux et al., 2008). The possibility of implicit change detection, or 'mindsight' (Simons and Ambinder, 2005), adds to the idea of illusion by suggesting that what people do is not always a consequence of what they consciously see, or report seeing.

Change blindness could have serious consequences in ordinary life, such as failing to detect changes while driving. Experiments have used natural traffic scenes with changes made during blinks, blanks, and saccades. Relevant changes are detected more quickly than irrelevant ones, but even so can take 180 ms longer to detect than when seen without any kind of disruption (Domhoeffler, Unema, and Velichkovsky, 2002). Not only do we blink and move our eyes while driving, but mud splashes on the windscreen can disrupt change detection, too. O'Regan, Rensink, and Clark (1999) showed that small shapes briefly spattered over a picture could prevent even large changes elsewhere being noticed. Comparable events happen all the time on the road and in the air, suggesting that dangerous mistakes might be made by drivers or pilots if a crucial event occurs just as some mud or a large insect splats onto the windscreen. Later experiments found, unexpectedly, that driving expertise made no difference to change blindness for driving-relevant changes, and also that relevant changes near the periphery of a driving scene were detected faster than those near the central vanishing point of the road ahead, where we would expect drivers to be focused (Galpin, Underwood, and Crundall, 2009).

These findings raise questions about how closely static images viewed on a screen can really replicate the conditions of real-life driving. One general criticism of these experiments is that the changes – like the height of a railing moving up a foot or two behind a couple's heads – are implausible, and there is no reason why a well-adapted visual system should be able to detect them. Some of the research on inattentional blindness does better at bringing the laboratory into the world.

INATTENTIONAL BLINDNESS

Could it be that if you don't pay attention to something, you simply do not see it? Just how paying attention to something relates to being conscious of it is hard to pin down (Chapter 7), but one way of thinking about the connection comes from studies of the odd phenomenon of inattentional blindness, pioneered by

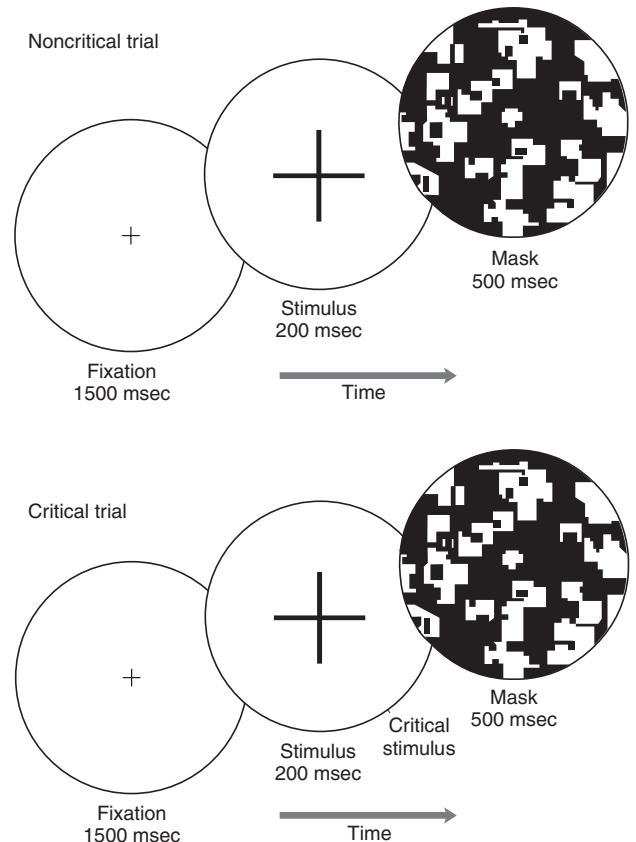


FIGURE 3.9 • Displays for the critical and non-critical trials in Mack and Rock's experiments. In this experiment the critical stimulus is in the parafovea. In other experiments the cross was in the parafovea and the critical stimulus was at the fixation point (Mack and Rock, 1998).

• SECTION ONE: THE PROBLEM

psychologist Arien Mack. On the basis of many experiments, he concluded that 'we rarely see what we are looking at unless our attention is directed to it' (2003, p. 180).

In a typical experiment participants are asked to look at a screen and fixate a marker (Mack and Rock, 1998). A cross briefly appears and they have to decide whether the horizontal or vertical arm is longer. Then on a critical trial, an unexpected stimulus appears – perhaps a black square or a coloured shape. Afterwards they are asked, 'Did you see anything on the screen on this trial that had not been there on previous trials?' On average 25% of participants say 'no'. This was true even when the cross they were attending to was slightly to one side of their fixation point and the unexpected shape appeared on their fovea. Indeed, they were even *less* likely to see the shape under these conditions (between 60% and 80% said they couldn't see it), suggesting that they had to actively *inhibit* attention at the fovea when trying to attend somewhere else. Interestingly, if the unexpected stimulus was a smiley face icon, or the person's own name, they were much more likely to notice it, suggesting that the unseen stimuli must still be processed to some extent.

A dramatic demonstration of inattentive blindness is the film starring 'gorillas in our midst' (Simons and Chabris, 1999). Two teams of students are seen throwing balls to each other, and observers are told to watch the white team very carefully and count the number of passes made. Afterwards they are asked whether they saw anything unusual in the film. What most usually miss is that a woman dressed in a gorilla suit walks right into shot, turns to camera and thumps her chest, and walks off on the opposite side. If you are an observer it is quite shocking to see the film again and realise what you missed. In experiments, approximately 50% of observers failed to notice the gorilla; they were more likely to see it when the counting task was easier or when they were watching the black team (since the gorilla was black).

In a NASA experiment something similarly striking – and somewhat more realistic – went unnoticed. Pilots were tested in a flight simulator and another plane was placed on the runway as they simulated landing. Novice pilots were more likely to see the obstruction, suggesting that the trained pilots failed to see something so very unexpected. This fits with results showing that the most important factor affecting inattentive blindness is a person's own attentional goals (Most et al., 2005). It is hard to see something that is truly unanticipated.

Attention may be necessary, but American psychologists Daniel Levin and Daniel Simons (1997) wondered whether it is sufficient (see also Chapter 7). In a change-blindness study they created short movies in which changes were made during a cut, either in arbitrary locations or at the centre of attention. In one movie two women are seen chatting over a meal. When the camera cuts to a new position one woman's scarf disappears, or the pink plates are changed to white. Few observers noticed these central changes. In another film an actor sitting at a desk gets up to answer the telephone and moves towards the door. The camera then cuts to a view in the hallway where a different actor answers the phone. When forty participants were asked to describe what they had seen, only 33% mentioned the changed actor. Apparently, attending to the main character in the film was not sufficient to detect the change. Levin and Simons conclude that even when we attend to an object, we may not form a rich representation that is preserved from one view to the next. So, although

a lack of attention in the right place might be able to account for many change-blindness results, paying attention to the right thing certainly does not guarantee that we will notice when it changes.

You might think that you are good at spotting the little mistakes made by TV producers and film directors, but these results suggest that very few people notice such inconsistencies – only those who happen to be attending to the detail in question, and sometimes not even then.

These effects are not confined to films and laboratory conditions. Simons and Levin (1998) arranged for an experimenter to approach a pedestrian on the campus of Cornell University and ask for directions. While they talked, two men rudely carried a door right between them. The first experimenter grabbed the back of the door and the person who had been carrying it let go and took over the conversation. Only half of the pedestrians noticed the substitution. Again, when people are asked whether they think they would detect such a change they are convinced that they would – but they are quite likely to be wrong. Change blindness and inattentional blindness are a great resource for magicians wishing to fool people.

More serious implications of inattentional blindness include, as with change blindness, effects on drivers and pilots. Talking on a mobile phone while driving is known to slow responses and increase errors but might also cause inattentional blindness. In experiments in which participants concentrated on tracking moving items in a dynamic display, a salient object suddenly appeared and was visible for 5 seconds (Scholl et al., 2003). Normally about 30% of people failed to see the unexpected event, but when they were simultaneously having a phone conversation, although their task performance did not suffer, 90% failed to detect it.

Later research showed that drivers' attention is impaired less when having a conversation with a passenger than when speaking to someone on the phone, presumably because the passenger is aware of the driving situation too. In a driving simulator involving a crowded motorway scenario, drivers talked either to a passenger or to someone using a handsfree phone, or used a handsfree phone enhanced by a video link showing their face and the driving scene. The last condition, which gave



CONCEPT 3.1 MAGIC

● If things that happen right in front of our eyes can be invisible, then magicians should be able to exploit this fact, and indeed, they have long done so. Now some, including James Randi and John Teller, are taking part in psychological research, too (Kuhn et al., 2008; Macknik et al., 2008). Strong magical experiences involve the illusion of impossibility and to create this, magicians must exploit our erroneous assumptions about the laws of nature or the huge gaps between what we see and what we believe we can see (Beth & Ekroll, 2015).

In a typical trick, the 'effect' is what the audience see (or think they see), and the 'method' is how the magician achieves the effect. For example, the audience may 'see' a coin passed from one hand to the other, when in fact it remained in one hand.

In physical misdirection, magicians use movement, high contrast, or surprise to direct interest and then carry out the method elsewhere. They can manipulate levels of attention with body language or joke and carry out the method as the audience relaxes. In psychological misdirection, they control expectation and surprise, give false clues to an impossible solution, or repeat the same effect using different methods. Above all, they manipulate the observer's gaze by knowing it will follow their own.

In a simple trick, the audience 'see' a ball fly into the air though it never left the magician's moving hand. Over a century ago, psychologist and magician Norman Triplett (1900) found that children 'saw' the ball disappear somewhere between the magician and the ceiling. In more recent research, 68% of observers claimed to see an imaginary ball when the magician's gaze followed it, compared with 32% when he watched his hand (Kuhn et al., 2008). Timing is critical, too: the illusion of a coin moving from one hand to another is weaker when the time interval between the false transfer and the opening of the second fist increases

(Beth and Ekroll, 2015). Other studies have investigated the importance of hand movement and object handling in false transfers (Otero-Millan et al., 2011; Phillips, Natter, and Egan, 2015), and the different responses of the brain to magic tricks as opposed to other surprising events (Parris et al., 2009).

The scientific study of magic, and the magical study of cognition, are expanding rapidly, and in some respects just catching up with what magicians have known for centuries. Researchers describe the development of 'a neurobiology of disbelief' (Parris et al., 2009) or 'a science centered around the experience of wonder' (Rensink and Kuhn, 2015) to tell us more about cognition, meta-cognition, and why our illusions sometimes give us such pleasure.

the driver's interlocutor a similar amount of information as a passenger typically has, reduced the number of collisions with merging vehicles to the same as the number that occurred when talking to a passenger in the adjacent seat (Gaspar et al., 2014). The researchers speculated that this is because the conversation partner cannot only warn the driver about unexpected events, but can also modulate their conversation depending on changing traffic, allowing the driver to devote more attention to driving and so reducing inattentive blindness.

It is also strikingly easy to induce 'blindness' during much simpler tasks, like walking. When people use mobiles while walking, they walk more slowly, change direction more frequently, are less likely to acknowledge other people, and are less likely to notice a unicycling clown ride past them (Hyman et al., 2010). So even with an activity as relatively undemanding as walking, we are highly susceptible to competing demands on our attention, to the point of effective 'blindness'. The

blindness can sometimes be useful, though, not only for preventing overload but also for allowing 'banner blindness' to unwanted advertising. Being selective is, after all, the blessing as well as the curse.

IMPLICATIONS FOR THEORIES OF VISUAL CONSCIOUSNESS

What do these results on change blindness and inattentive blindness mean for consciousness? They seem to challenge the simplistic 'stream of vision' theory, and to imply that vision is not a process of building up detailed inner representations that can be used to compare the scene from one moment to the next. They suggest that we do not hold on to nearly as much information as we seem to need, and 'the richness of our visual world is, to this extent, an illusion' (Blackmore et al., 1995, p. 1075). Yet obviously something is retained, otherwise we would have no sense of continuity and not even notice if the entire scene changed. So we should be cautious about leaping to exaggerated conclusions (Simons and Rensink, 2005).

The results do not prove that we never have a detailed representation of the scene, or that during a saccade we retain no representation of what was seen just before. We might fail to detect a change even with representations of both the pre-change and post-change scenes if a very short-lived initial representation were overwritten by the next scene. We might have detailed representations of both versions of the scene available at once but fail to compare them properly. Or we might accurately represent details of the original scene but then not update them when the scene changes, believing we have already extracted all the meaning we need. We may even cleverly combine both representations, retaining some features from each and so never notice that there ever were two versions. There are therefore many possible interpretations, and theorists vary in how much, and

what sort of, information they claim is retained, for how long, and what is done with it (Simons, 2000).

Those who emphasise the lack of detail retained in the internal representation might describe it in terms of a 'sketchy higher-level representation' (Blackmore et al., 1995) or 'extremely reduced visual representations' (Hayhoe, 2000). Two other possible ways of thinking about it are as a 'gist' and a 'virtual representation'.

The idea of the gist was proposed as part of a straightforward interpretation of change blindness initially given by Simons and Levin (1997). During any single fixation we have a rich visual experience. From that, we extract the meaning or gist of the scene. Then, when we move our eyes, we get a new visual experience, but if the gist remains the same our perceptual system assumes the details are the same, and so we do not notice changes. This, they argue, makes sense in the rapidly changing and complex world we live in. We get a phenomenal experience of continuity without too much confusion.

Somewhat more radically, Canadian psychologist Ronald Rensink (2000) suggests that observers never form a complete representation of the world around them – not even during fixations – and that there is no visual buffer zone which accumulates an internal picture. Instead, object representations are built one at a time, as needed. Focused attention takes a few proto-objects from low-level processing and binds them into a 'coherence field' representing an individual object that persists for some time. When attention is released the object loses its coherence and dissolves, or falls back into an unbound soup of separate features.

To explain why we seem to experience so many objects at once, when so little is held in focused attention, Rensink argues that vision is based on 'virtual representations': these are constructed from gist, spatial layout, and a longer-term schema of the scene. They are not 'structures *built up* from eye movements and attentional shifts, but rather, are structures that *guide* such activities' (2000, p. 36). We get the impression of a rich visual world because a new representation can always be made 'just in time' using information from the world itself. Sometimes such representations may be stable; sometimes they may contain a large amount of detail. But at no point are they both stable and detailed. Indeed, eye-tracking evidence suggests that representations of the world may last only about the same length of time as a fixation between eye movements (Tatler and Land, 2011).

O'Regan (1992) agrees that we do not need to store large amounts of visual information because we can use 'the world as an external memory', or as its own best model, but he goes so far as to reject the idea that we need to make our own internal models at all. He criticises traditional theories of vision for being based (even if they don't admit it) on the assumption that in visual perception, the distortions and gappiness of the retinal image are compensated for by the brain's construction of a detailed model of the outside world, which somehow creates perceptual consciousness. Instead, he argues that the visual world is not something that we have, or build up, but something we do.

O'Regan and Noë (2001) proposed a sensorimotor theory of vision and visual consciousness – a new way of thinking about vision that owes a debt to thinkers like the phenomenologist Merleau-Ponty, who considered the body in action

'we can use the world as its own best model'

(Clark, 1997, p. 29)

*'seeing is a way of acting
[...] of exploring the
environment'*

(O'Regan and Noë, 2001,
p. 939)

*'sensations are
representations of
something we do'*

(Humphrey, 2016, p. 117)

essential to understanding consciousness: 'Consciousness is in the first place not a matter of "I think that" but of "I can"' (1945/2002, p. 159). O'Regan and Noë argue that classical theories of vision do not explain how the existence of an internal representation can give rise to visual consciousness (another version of the hard problem). In their theory, the hard problem is avoided because 'The outside world serves as its own, external, representation'. Instead of being about building representational models, 'seeing is a way of acting. It is a particular way of exploring the environment' (2001, p. 939).

More specifically, an organism has the experience of seeing when it masters the governing laws of sensorimotor contingencies – that is, when it develops skills for extracting information from the world; for interacting with the visual input and exploiting the ways in which that input changes with eye movements, body movements, blinks, and other actions. On this view what you see is those aspects of the scene that you are currently 'visually manipulating'. If you do not interact with the world, you see nothing. When you stop manipulating some aspect of the world, it drops back into nothingness.

As with Rensink's virtual representation, what remains between saccades is not a picture of the world, but the information needed for further exploration. A study by Karn and Hayhoe (2000) confirmed that spatial information required to control eye movements is retained across saccades. Could this be sufficient to give an illusion of continuity and stability?

This theory is radically counter-intuitive, not least because seeing does not feel like manipulating temporary aspects of the world that then disappear, but O'Regan likens it to the light inside your fridge. Every time you open the door the light is on. Then you close it and wonder whether it's still on. So you open it and look again. It's still on. So it is with the world: it is always there when you look, so it's easy to think you have a constant detailed representation of it.

Sensorimotor theory is dramatically different from most existing theories of perception, but is closely related to theories of embodied or enactive cognition (Chapter 8). It is similar to the idea of perception as a kind of 'reaching out' (Humphrey, 1992, 2006), to theories stressing the interdependence of perception and action (Hurley, 1998), to J.J. Gibson's (1979) ecological approach to perception, and further back to Merleau-Ponty's idea that 'consciousness is nothing other than the dialectic of milieu and action' (1942/1965, pp. 168–169). Seeing does not mean building representations of the world that can then be acted upon; rather, seeing, attending, and experiencing are all kinds of action.

SEEING OR BLIND? A THOUGHT EXPERIMENT TO TEST SENSORIMOTOR THEORY

According to O'Regan and Noë's (2001, p. 1020) sensorimotor theory, 'perception is constituted by mastery of sensorimotor contingencies'. Seeing means manipulating the contingencies between action and input, such as moving one's eyes and getting changed visual input. A thought experiment suggests a bizarre consequence of this theory.



FIGURE 3.10A • Alva at start

CONCEPT 3.2

Thinking about visual consciousness as a way of acting rather than a stream of pictures exposes some deep-seated patterns in how we think about our experience. These theories run counter not only to more traditional theories of vision, but also, as we saw earlier, to our intuitive understanding of how vision works. This might make us wonder: how on earth could I be so mistaken about my own consciousness? But there are good reasons why you might be. Whatever the details of whether and how it constructs neural models of the world, the visual system is remarkably fit for purpose: its complex parallel processes of texture recognition, pop-out detection, contrast control, and all the rest work so well together that it is easy to believe that the parts add up to a detailed picture-like whole. But the picture may be only the result we infer, not in fact the mechanism we rely on. Could we learn to see this inference emerging, and stop it in its tracks by reminding ourselves that the picture may be an effect rather than the cause of our ability to see so well?

The idea of doing this may be a little scary. The appeal of the ‘picture in the head’ model lies partly in the stability and security it promises; how could we function safely and reliably in the world without such a picture? Don’t we need an accurate image of the world, transmitted from eyes to brain and then available to all the processes directing the rest of cognition and motor action, to think and act appropriately? The appeal is also tightly tied to convention: it is hard if not impossible to separate how things seem visually from the habits for seeing that we learn from society and culture. The invention of the camera changed how we see the world (Berger, 1972) and our world is more and more dominated by images designed to capture and hold our attention. Are they making the ancient intuitions about pictures in the head ever harder to dislodge?

We began with the idea of a stream of vision, and the assumption that it is a stream of internal pictures or representations. The results on filling-in, inattention blindness, and change blindness all call that idea into question (Blackmore, 2002), raising the possibility that vision may be a grand illusion. This is just one aspect of the illusionist proposal (e.g. Frankish, 2016a) that our ideas about all of our experience are wrong.



FIGURE 3.10B • Kevin option 1



FIGURE 3.10C • Kevin option 2

One participant, Kevin, wears a head-mounted display which shows the output from an eye tracker on a second person, Alva. When Alva moves or looks around the room, everything he looks at is instantly fed to ‘Kevin’. Kevin therefore gets exactly the same visual input as Alva. Kevin is also making eye movements, but although Alva’s eye movements correspond to the changes of scene, Kevin’s do not. This means that Kevin can have no mastery of sensorimotor contingencies when moving his eyes around.

What will happen? You might like to consider your own answer before reading on. Here are three possibilities:

- 1 Kevin can see perfectly well. He is receiving the same visual input as Alva and so must see the same as Alva sees.
- 2 Kevin is effectively blind because although he receives the same visual input as Alva he cannot master the contingencies between input and eye movements.
- 3 Perhaps Kevin can see something but not the same as Alva.

In her peer commentary on the 2001 paper, Sue (Blackmore, 2001) suggested that the sensorimotor theory makes the strong prediction that Kevin is effectively blind and unable to recognise things, judge distances, grasp objects, or avoid obstacles. Possibly he might have some other residual vision, but if eye movements were uncorrelated with input the mainstay of what it means to see would be gone.

In a poster (Blackmore, 2007a), she asked participants at a vision conference, including O’Regan and Noë, to give

their opinions, and gathered others online. The results are shown in the table. Those who chose Outcome 1 are effectively rejecting O'Regan and Noë's theory even if they say they agree with it. Those who choose outcome 2 are making the strong – indeed, extraordinary – prediction that it is possible for two normally sighted people who receive identical visual input to have completely different experiences. If true, this would suggest that our illusions about vision are further-reaching than we thought. You can read Sue's commentary on p. 977 of the 2001 paper, and the authors' response on p. 1020.

Sensorimotor theory			
	1 Can see	2 Is blind	3 Other
True	5	11	6
False	6	3	5

FIGURE 3.10D • Responses to the poster (Blackmore, 2007a). As expected, the majority of those who think sensorimotor theory is true think Kevin must be blind, and those who do not think he can see. Yet there are some who think it is true and still think Kevin can see. This is only a thought experiment but may be able to help us think about the consequences of this counterintuitive theory.

If we were convinced that the grand illusion theory was correct, we might perhaps begin to experience vision as a form of acting in the world, and maybe then seeing would cease to seem like a stream of pictures. In this way, escaping the illusion could really change the way we see the world.

This possibility may be open to us in other areas too. Other elements of our conscious experience, which are equally crucial to our sense of who we are, may also be subject to illusions. Later in this book we will tackle the issue of free will (Chapter 9), which many people consider essential to be what makes us human, and to be what makes us good (not evil) and responsible (not apathetic). We will see that there are strong arguments for thinking about free will too as an illusion: something which is not what it seems. And we hope that as we work our way through these ideas, a thread running through the whole book will become increasingly clear: that the very construct everything seems to pivot around, the idea of the conscious me, myself, my *self*, might be illusory, too (Chapter 16).

Meanwhile, however, we will conclude with one last thought about illusions. The very idea that it might be possible to be mistaken about our own consciousness is a tricky one. If we try to distinguish sharply between 'consciousness itself' and 'how consciousness seems to us', we end up believing there are two separate things to explain and then realising

that each must have effects on the other: how we think and talk about our conscious experience inevitably affects that experience. So maybe the very idea of an illusion is itself mistaken because it requires there to be a reality of conscious experience that we can be mistaken about. Maybe this is not so and, as many illusionists believe, our ways of being mistaken are all there is.

HOW MUCH AM I SEEING NOW?

'We can be Realists about qualia, or else we have to be Illusionists'

(Humphrey, 2017)

'Illusionists deny that experiences have phenomenal properties, and focus on explaining why they seem to have them'

(Frankish, 2016b, p. 14)



FIGURE 3.11 • Is the light always on inside the fridge?



Dennett, D. C. (1991). Dismantling the witness protection program [excerpt]. *Consciousness explained* (pp. 344–356). London: Little, Brown.

Presents the view that filling-in is unnecessary.

Komatsu, H. (2006). The neural mechanisms of perceptual filling-in. *Nature Reviews Neuroscience*, 7, 220–231.

Provides evidence from neuroscience that filling-in does occur.

Noë, A. (Ed.) (2002). Is the visual world a grand illusion? Special issue of the *Journal of Consciousness Studies*, 9(5–6).

A brief preface by Noë is followed by fourteen articles debating the grand illusion.

O’Regan, J. K., and Noë, A. (2001). A sensorimotor account of vision and visual consciousness. *Behavioral and Brain Sciences*, 24(5), 883–1031 (incl. commentaries and authors’ response).

A radical break with the idea that vision depends on building up representations of what is being seen.

Simons, D. J., and Rensink, R. A. (2005). Change blindness: Past, present, and future. *Trends in Cognitive Sciences*, 9(1), 16–20.

What can and can’t we conclude from change blindness, and what further questions does it raise?

Rensink, R. A., and Kuhn, G. (2015). A framework for using magic to study the mind. *Frontiers in Psychology*, 5, article 1508.

Explores how magicians’ methods and effects may contribute to a psychology of consciousness and the experience of wonder.



The brain
SECTION
Two



Neuroscience and the correlates of consciousness

CHAPTER

FOUR

If you could look right inside a brain and see everything that was happening there, would you then understand consciousness?

Identity theorists say yes: the mind and the brain are identical. If we could observe brain activity in sufficient detail and at many different levels of organisation, then we would understand everything that the brain was doing, and since consciousness *is* the activity of brains, it follows that we would understand consciousness. As philosopher Dan Lloyd puts it, we would find 'that there is in fact just one system, and that the neural version and the phenomenal version are simply different labels applied to one underlying Reality' (2004, p. 299).

'The entire brain is sufficient for consciousness'

(Koch, 2004, p. 87)

Some eliminative materialists say yes, too. By definition, eliminative materialists *eliminate* mental properties like qualia: they claim that the mental states we assume to exist actually do not. Although there is no reason why eliminative materialists should necessarily think that the brain must provide the whole solution, many do, arguing that what the mind does is nothing more than what the brain does.

'Extended minders' and other theorists of embodied cognition say no. They insist that neural activity alone cannot provide any answers, and that we need to take the rest of the body and the environment into account too. 'You are not your brain', says Alva Noë. 'Consciousness does not happen in the brain. That's why we have been unable to come up with a good explanation of its neural basis' (Noë, 2009, p. 5). We must include the person's history, the world around them, and the whole body's interactions with that world. From this perspective, the mistake of neurocentrism is a form of the mereological fallacy of ascribing to part of an

'Consciousness does not happen in the brain'

(Noë, 2009, p. 5)

SECTION TWO: THE BRAIN




FIGURE 4.1 • Schematic illustrations of the human brain. (a) A lateral view (i.e., looking at the outside of one side of the brain) of the left hemisphere showing the four lobes of the cortex and the various sensory and association areas. (b) A medial view (i.e., looking at the inside surface of one half of the brain as though it has been cut through the middle) of the right hemisphere. The corpus callosum consists of over 200 million nerve axons connecting the two hemispheres. Thalamic and midbrain structures are also shown.

animal 'an attribute which it makes sense to ascribe only to the animal as a whole' (Bennett and Hacker, 2003, p. 240). British philosopher Andy Clark (Clark and Chalmers, 1998; Clark, 2008) conceives of a person as an extended or 'supersized' system whose 'operations are realized not in the neural system alone but in the whole embodied system located in the world' (Clark, 2008, p. 14). (We will return to these views in Chapter 8.)

Mysterians also say no, but for a different reason. Many of them claim that we can never understand consciousness: it is simply not something the human mind is capable of grasping. Some, like Steven Pinker (2007, p. 6), admit that we might one day be able to, but think it is pretty unlikely:

The brain is a product of evolution, and just as animal brains have their limitations, we have ours. Our brains can't hold a hundred numbers in memory, can't visualize seven-dimensional space and perhaps can't intuitively grasp why neural information processing observed from the outside should give rise to subjective experience on the inside. This is where

I place my bet, though I admit that the theory could be demolished when an unborn genius – a Darwin or Einstein of consciousness – comes up with a flabbergasting new idea that suddenly makes it all clear to us.

No one denies that the brain is relevant to consciousness; they just disagree fundamentally about its role. Looking inside a brain reveals a mystery whichever method you use. Dissecting a human brain with a scalpel and looking with the naked eye reveals a few pounds of soft greyish tissue with a wrinkly surface and not much inner detail. Staining a slice of brain and looking through a microscope shows billions of neurons with vast spreading trees of axons and dendrites. Attaching electrodes to the scalp provides a readout of activity on the surface,

and modern methods of scanning and statistical analysis give multicoloured representations of what is happening inside. But in every case the mystery is obvious – how can this physical lump of stuff, with its electrical and chemical activity, relate to conscious experience? Whatever the answer, it is worth learning a little about the structure and function of the human brain.

A HUMAN BRAIN

Said to be the most complex object in the known universe, a human brain contains about 86 billion neurons connected by trillions of synapses between them, along with billions of supporting glial cells, some of which are also involved in signalling. Human brains are much larger, relative to body weight, than those of any other animal, but are organised in a roughly similar way. Sensory and motor neurons from all over the body run into the spinal cord and up into the brainstem at the base of the brain. All these neurons are collectively known as the peripheral nervous system, while the spinal cord and brain make up the central nervous system (CNS).

The brainstem, consisting of medulla, pons, and mid-brain, is essential for life, not only because it carries so many important nerve tracts but because of its role in controlling cardiac, respiratory, and sexual functions and arousal levels.

The reticular formation in vertebrates is involved in the pain desensitisation pathway and, along with its connections, forms the reticular activating system, which activates widespread regions of the cortex in transitions from sleep to waking or from relaxed waking to alert attention. It has been known since the nineteenth century that animals with no cortex can still show normal sleep–waking cycles controlled by this system. Its functioning is thought necessary, but not sufficient, for consciousness.

Behind the midbrain is the cerebellum, or ‘little brain’, whose main function is motor control, with extensive links upwards to motor cortex and downwards through the spinocerebellar tract, which provides feedback on body position and the effects of actions.



CONCEPT 4.1 MAPPING THE BRAIN

• Single cell recording

Fine electrodes are inserted into living cells to record their electrical activity. This technique is widely used in animal studies, and more rarely in humans.

Electroencephalogram (EEG)

The EEG uses electrodes on the scalp to measure changes in electrical potential arising from the combined activity of many cells in the underlying area of the brain. The human EEG was first described in 1929 by the German psychiatrist Hans Berger, who showed that the resting alpha rhythm (8–12 cycles per second) is blocked by opening the eyes or doing mental arithmetic. In the 1960s, computer averaging improved the study of event-related potentials (ERPs), including evoked potentials in response to specific stimuli, readiness potentials

that build up gradually before a response is made, and potentials associated with unexpected events. Although the EEG has poor spatial resolution, it is still a valuable research tool because of its good temporal resolution.

X-ray Computed Tomography (CT)

Developed in the early 1970s, CT scans are computer-generated images of tissue density, produced by passing X-rays through the body at many different angles and measuring their attenuation by different tissues. The same mathematical techniques for constructing the images are used in newer forms of scanning.

Positron Emission Tomography (PET)

This is a technique for imaging the distribution of radioactivity following administration of a radioactive substance. In PET, atoms that emit positrons are incorporated into glucose or oxygen molecules, allowing brain metabolism and blood flow to be measured directly. Radiation detectors are arrayed on the head in several rings, allowing several slices of brain to be studied simultaneously. PET has good spatial resolution but very poor temporal resolution, and the added disadvantage of having to use radiation.

Nuclear Magnetic Resonance (MRI)

MRI measures the radio signals emitted by some atomic nuclei (e.g. ^1H , ^{13}C , and ^{31}P) when placed in a magnetic field and excited by radio frequency energy. The radiation emitted provides information about the chemical environment of the nuclei. In the 1970s, the idea of using hydrogen atoms in the body for imaging was developed into fMRI (functional MRI), which can provide extremely detailed images of living brains. Early methods required injections of a paramagnetic substance, but in the 1990s totally non-invasive methods followed, including the use of BOLD (blood oxygen level-dependent) contrast, which allows measurement of local brain metabolism. fMRI measures neuronal activity only indirectly, depending on metabolic and haemodynamic responses to neural activity, which limits its temporal resolution. For brain scanning, the head has to be placed inside the scanner and kept very still. The results are displayed using false colour to produce the familiar coloured images of the brain in action. Although they may look like direct representations of brain activity, the published images have gone through many stages of processing and statistical analysis and must be interpreted with care: the readings are subject to noise at every stage of the process; false positives are extremely easy to generate when a set of 'standard assumptions' are not met (Eklund, Nichols, and Knutsson, 2016) – so easy that a dead Arctic salmon can appear to be engaged in a perspective-taking task (Bennett, Miller, and Wolford, 2009); and basic variables like breathing may be serious confounds (Birn et al., 2006; Huijbers et al., 2014). There are now numerous open-science initiatives for enhancing validity and reproducibility in fMRI research, but we must bear in mind its limitations.

Transcranial magnetic stimulation (TMS)

In TMS, or repetitive TMS (rTMS), a coil held over the brain generates a pulsed magnetic field which stimulates neurons in a focused area by inducing small local currents. Stimulating motor areas induces involuntary movements, and if the precise area stimulated is located by scanning, this allows motor cortex to be accurately mapped. Similarly, visual or speech areas can be mapped because TMS suppresses function in the area stimulated. TMS can also be used to induce particular experiences or altered states of consciousness (Chapter 13).

Between the midbrain and the cortex is the thalamus, which contains relay areas for sensory inputs including vision, hearing, and touch, as well as motor functions. These 'relays' do not just pass signals on but form crucial parts of complex loops to and from the cortex which lies above and around it. These thalamocortical loops have been implicated in consciousness (e.g. Llinas and Ribary, 2001; Trapp, Schroll and Hamker, 2012) and are especially well developed in humans.

Finally, there is the cortex, or outermost layer of the brain. Its oldest and innermost part, the limbic system, is common to many other animals and is sometimes referred to as the reptilian brain. This includes many structures implicated in consciousness: the hippocampus, essential for laying down long-term memories and forming cognitive maps; the amygdala, with roles in rewards and emotions; the hypothalamus, which regulates the autonomic system including blood pressure, heart rate, and sexual arousal; and the cingulate gyrus, which is involved in emotion, pain, and motivational responses. These are all hidden underneath the neocortex.

A man who attempts to ease the miseries of failing minds by repairing brains is bound to respect the material world, its limits, and what it can sustain – consciousness, no less. It isn't an article of faith with him, he knows it for a quotidian fact, the mind is what the brain, mere matter, performs. If that's worthy of awe, it also deserves curiosity; the actual, not the magical, should be the challenge.

(Ian McEwan, *Saturday*, 2005, p. 67)

The neocortex has expanded more than any other part during human evolution, becoming deeply folded to give a large surface area of grey matter (neuron cell bodies and unmyelinated, or electrically uninsulated, fibres) on top of the white matter (myelinated axons). Its two main types of neurons are excitatory pyramidal cells and inhibitory interneurons. Most of the cortex is arranged in six layers, with layer 1 on the outside. There are also vertical columns showing functional organisation, such as

those relating to neighbouring areas of skin or muscles, or the visual field. Sensory areas are organised in a roughly hierarchical manner: levels of processing build on each other and neurons doing similar jobs are close together; but no area is isolated from the rest and everywhere there are long cortico-cortical connections and cortico-thalamic loops providing a massively interlinked system with no ultimate top.

The two hemispheres of the cortex are linked by the white matter of the smaller anterior commissure and much larger corpus callosum, which is a wide band of about 200 million fibres beneath the cortex. Each hemisphere has four lobes (see Figure 4.1). Although these were originally labelled by appearance, they turn out to be roughly divided by function: the occipital lobe deals with vision; the parietal lobe includes sensory association areas as well as somatosensory cortex and the dorsal stream; the temporal lobes include auditory areas and memory functions as well as the ventral stream; and the frontal lobes, which are especially large in humans, deal with forward planning and executive functions.

To understand the neural basis of consciousness, far more detail is needed (see e.g. Baars and Gage, 2010), but this superficial overview should be enough of a guide for those whose primary interest is psychological or philosophical. We can now begin that look inside the brain. But what do we look for?

The most popular method of evading the mystery (or tackling the 'easy problems') is to look for 'the neural correlates of consciousness' (NCCs).

CORRELATIONS BEFORE CAUSE

The idea behind studying the NCCs is to measure some aspect of neural functioning and then correlate it with reports of conscious experience (Metzinger, 2000). The 'contrastive method' involves comparing measurements of neural functioning when a given action or perception is conscious with when it is not conscious (Baars, 1997a, 1997b; Aru et al., 2012). But which aspect of neural functioning? Measurements have been made, and theories proposed, at every scale from single molecules to large-scale assemblies of neurons, for it is not yet obvious what we should be looking for.

The classic method for looking for the NCCs is to focus on vision, and in particular the phenomenon of rivalry, in which perception alternates between different options.

One form of rivalry involves ambiguous figures. For example, look at the Necker cube in Figure 4.3. Keep your eyes fixated on the central dot and watch what happens. This simple figure can be interpreted in depth in two mutually incompatible ways. Even though you are keeping your eyes still, you should find that the cube flips back and forth between the two different interpretations. It is impossible to see both at once, or to combine the views into one, so you experience alternation or rivalry between the two.




FIGURE 4.2 • Some of the basic types of neuron




FIGURE 4.3 • The Necker cube: a simple example of rivalry. Keep your gaze on the central spot while looking at the cube. There are two equally likely interpretations that tend to alternate: one with the front face up and to the left, the other with the front face down and to the right. You may be able to flip views deliberately and vary the speed of alternation.

Binocular rivalry, by contrast, is found when different images are presented to the two eyes. For example, a picture of the seaside might be shown to the right eye and a face to the left, or a vertical grating to the left eye and a horizontal grating to the right. In such cases the face and ocean are not combined into one picture, nor do the gratings fuse into a plaid. Instead perception seems to flip between the two.

What is going on? Early theories suggested that the flipping was due to eye movements or other peripheral effects, but keeping the eyes still does not stop the alternation, and peripheral theories have not generally fared well. It seems more likely that the flipping occurs further up the visual system. But how does this relate to the subjective experience? It feels as though the two views are competing for consciousness. It seems as though first one, and then the other, gains access to consciousness and thus you become aware of it. This simple phenomenon provides an ideal situation for investigating the relationship between the objective facts (input to the eye, events in the visual system, and so on), and the subjective facts (being conscious of first one of the pictures and then the other).

The first experiments to use binocular rivalry to look for the NCCs were done with macaque monkeys (Logothetis and Schall, 1989; Sheinberg and Logothetis, 1997). Macaques can be trained to report which of two pictures they are seeing by pressing a lever, and their responses are much like ours. For example, when shown a vertical grating to one eye and a horizontal grating to the other, or gratings moving in different directions, they can press a lever to indicate when what they see flips from one to the other. Logothetis and his colleagues trained monkeys in this way, recording from single cells in various brain regions. They were looking for areas where the activity corresponded not to the unchanging visual input, but to the changing perceptions reported by the monkey's behaviour.




FIGURE 4.4 • The principle behind Logothetis's experiments. When monkeys are shown a different display to each eye they report binocular rivalry just as humans do. They cannot speak, but they can indicate which display they are currently seeing by pressing a lever.

Cells in early visual cortex, such as area V1, responded to the unchanging input. For example, some cells responded to vertical stripes, some to motion in different directions, and some to particular stimuli, but their behaviour did not change when the monkey's perception changed. Further along the visual pathway (for example in MT and V4), some cells responded to what the monkey reported seeing. Finally, in the inferior temporal cortex (IT), almost all the cells changed their response according to what the monkey reported seeing. So, if the monkey pressed the lever to indicate a flip, active cells stopped firing and a different set started. It looked as though activity in this area corresponded to what the monkey was consciously seeing.

Does this mean that the NCC lies in IT? One problem is that the connection with consciousness depends on assuming that the monkeys are consciously perceiving. This seems reasonable given the way they respond, but of course we cannot know for sure, and those who believe that language is necessary for consciousness might argue that the monkey's responses tell us nothing about human consciousness (Chapter 10). Since these early experiments, technology has made similar experiments on humans possible, and we

will come back to these later in this chapter. For now, we suggest that you keep this basic experimental method in mind as we consider some complex questions raised by this research.

Some neuroscientists, especially Francis Crick and Christof Koch (2003), have argued for ignoring the theoretical and philosophical problems and getting on with the search for correlations across a wide range of conscious experiences.

I argue for a research program whose supreme aim is to discover the neuronal correlates of consciousness, the NCC. These are the smallest set of brain mechanisms and events sufficient for some specific conscious feeling, as elemental as the color red or as complex as the sensual, mysterious, and primeval sensation evoked when looking at the jungle scene on the book jacket.

(Koch, 2004, pp. xv–xvi)

When thinking about NCCs it is important to remember the usual warnings about the meaning of ‘correlation’, above all that a correlation does not imply a cause. This familiar trap is especially easy to fall into when dealing with something as slippery as consciousness. When any correlation is observed between two events, A and B, there are three possible causal explanations: A caused B; B caused A; or some other event or process, C, caused them both. Alternatively, A and B might actually be the same thing even though they do not appear to be.

Correlation is not causation

In some cases, the right explanation is obvious. Imagine that you are at a railway station and every so often you see hundreds of people gathering on the platform, always followed by a train arriving. If correlation necessarily implied cause you would have to conclude that the people on the platform caused the train to appear. Obviously you won’t, because you know that both events were caused by something else: a railway timetable. When it comes to consciousness, however, things are not that obvious, and we can easily jump to false conclusions. According to Dan Wegner’s (e.g. 2005) theory of conscious will (Chapter 9), it is precisely this kind of confusion between correlation and cause which creates the illusion that our thoughts cause our actions, when in fact both are caused by prior neural activity. Similarly, we saw in Chapter 3 that rather than pictures in the head creating a rich and unified visual experience, it may be that the illusions of picture-like experiences and picture-like mechanisms are both caused by the adaptive fitness of our visual and motor systems.

So, when correlations are found between neural events and conscious experiences, we must consider all these possibilities. Perhaps neural events cause conscious experiences. Perhaps conscious experiences cause neural events. Perhaps something else causes both of them. Perhaps neural events *are* conscious experiences. Perhaps we have so misconstrued one or the other that none of these is true.

CONSCIOUS VISION

Francis Crick, famous as the co-discoverer of the structure of DNA, was among the first to look for the NCCs. Dismissing philosophy and adopting a thorough-going reductionist approach, he argued that ‘that we shouldn’t approach the hard

'the minimal neuronal events and mechanisms jointly sufficient for a specific conscious percept'

(Koch, 2004, p. 104)

problem head on. We should try and find the neural correlates that correspond to what we're conscious of' (in Blackmore, 2005, p. 69). He explained that he is looking for the correlates of the 'vivid representation in our brains of the scene directly before us' (Crick, 1994, p. 207). Until his death in 2004, he worked closely with Christof Koch to find 'the minimal neuronal events jointly sufficient for a specific conscious percept' (Koch, 2004, p. 104). So, rather than looking for the NCs of consciousness in general, they were looking for NCs of particular experiences. Crick says that they chose vision 'because humans are very visual animals and our visual awareness is especially vivid and rich in information' (Crick, 1997, p. 21). Also, visual inputs are relatively easy to control, we have detailed knowledge of the primate visual system, and that of higher primates is similar to our own. In some ways it is regrettable that vision has been studied in so much more depth than other senses, but it has a crucial place in the search for the NCCs. So, in this section we extend the discussion of the previous chapter to delve further into the neuroscience of vision. Note, however, that much research in this area assumes rich mental representations underlying visual experience, an assumption which Chapter 3 showed to be questionable.

'so far we can locate no single region in which the neural activity corresponds exactly to the vivid picture of the world we see in front of our eyes'

(Crick, 1994, p. 159)

At the start of their endeavour Crick said, 'so far we can locate no single region in which the neural activity corresponds exactly to the vivid picture of the world we see in front of our eyes' (1994, p. 159), but he knew what he was looking for – something which corresponds to that 'vivid picture'. He and Koch laid out their working hypotheses as a 'framework for consciousness' (2003). They proposed the front of the brain as a kind of unconscious homunculus observing the sensory areas, with many 'zombie' modes of processing all over the brain, consisting of transient coalitions of neurons corresponding to representations of thoughts, images, and perceptions. This idea of coalitions or neural assemblies goes back more than half a century to Donald Hebb (1949), but has been transformed by a better understanding of how large collections of neurons can work together. Crick and Koch proposed that these constantly changing coalitions compete with each other, attention biasing their competition. In vision, the neural activity travels fast up the hierarchy to frontal cortex, providing a conscious gist of the scene, and then travels more slowly back down the hierarchy to fill in the details. Recalling the picture-in-the-head theories discussed in Chapter 3, Crick and Koch proposed that conscious vision is like a series of snapshots with motion 'painted' on.

With this framework in place, they tried to find the NCCs. 'First you want an idea of whether it's that set of cells firing, or whether they fire in a special way, or whether it's a combination of the two, or something else quite different' (in Blackmore, 2005, p. 70). Crick was referring here to the different possible ways of thinking about the NCCs – as a place, as a specific group of neurons, or as a particular pattern of cell firing. The problem here is that if some processing is conscious and some is not, what is the 'magic difference'? Do some cells have a special extra ingredient? Are some patterns of firing able to 'create' or 'give rise to' subjective experiences, while others cannot? Does connecting cells up in a special way, or in certain sized groups, make consciousness happen in those cells but not in others?

Put like this, none of the options sound very plausible. If you're looking for the NCCs, by definition you believe in the hard problem, but it often seems that instead of solving it you merely end up suggesting that it applies to some brain areas or processes and not others. Yet this does not deter researchers from

describing the areas identified in the neuroimaging studies as visual consciousness areas (ffytche, 2000) or sites where consciousness is generated (Chalmers, 2000), and the brain processes involved as those 'that are qualia laden as opposed to those that are not' (Ramachandran and Hubbard, 2001, p. 24). Meanwhile, Crick and Koch suggest that 'it is the transient results of the computations which correlate with qualia; most of the computations leading up to those results are likely to be *unconscious*' (Crick and Koch, 2000, p. 104, original italics). Similarly, Ray Jackendoff (1987) and Jesse Prinz (2007) try to determine what kind of information is encoded in different stages of the visual hierarchy and how well it corresponds to the features of conscious visual experience. Prinz asks 'Where, in the flow of information, does consciousness arise?', and concludes that 'Conscious states are attentionally modulated intermediate-level representations' (pp. 247, 258).

The idea of a place where consciousness happens is what Dennett characterises using the metaphor of the Cartesian theatre (which we met in Chapter 1 and will learn more about in the next chapter): the place where everything comes together and consciousness happens. It is seen in extreme form in Descartes's idea of the pineal gland as the seat of the soul, and in the view that William James pilloried of a single 'pontifical neuron' to which 'our' consciousness is attached. We know that damage to almost any area of the brain has some effect on consciousness, and so in some sense the whole brain is involved. This was certainly James's view. He said that 'The consciousness, which is itself an integral thing not made of parts, "corresponds" to the entire activity of the brain, whatever that may be, at the moment' (1890, i, p. 177). But he was under no illusion that this solved the problems: 'The ultimate of ultimate problems, of course, in the study of the relations of thought and brain, is to understand why and how such disparate things are connected at all' (1890, i, p. 177). This, his version of the hard problem, remains whichever areas are favoured, and whether they are discrete or distributed.

The idea of a 'bridge locus' was proposed by vision researchers in the 1980s as a locatable part of the brain to be a crossing point for the chasm between inner and outer worlds, between matter and consciousness (Movshon, 2013). Neural activity at the bridge locus would be necessary and sufficient for visual consciousness: 'The occurrence of a particular activity pattern in these bridge locus neurons is necessary for the occurrence of a particular perceptual state; neural activity elsewhere in the visual system is not necessary' (Teller and Pugh, 1983, p. 581). Much neuroscientific work on visual consciousness, even if it has not explicitly invoked the bridge locus, seems to assume that something like it might be found. As vision scientist Anthony Movshon puts it: 'In neuroscience, a huge premium is now placed on localizing particular functions in the brain. No conversation about perception, cognition, or action goes far without a consideration of where in the brain the function under discussion is "located"' (2013, p. 221).

We must remember here the difference between identifying brain areas involved in specific cognitive functions and trying to find those responsible for conscious experiences. While fMRI and other scanning methods find ever more locations where activity correlates with particular experiences or actions, this is very different from research on the NCCs, which tries to find those specifically responsible for conscious experiences as opposed to unconscious processing.

We can see this difference with respect to vision. The visual system is well understood, with something like ten separate parallel pathways from the eyes to different areas of the brain. About 85% of cells take the major route through the lateral geniculate nucleus (LGN) in the thalamus to primary visual cortex (V1) in the occipital lobe and then, with increasing numbers of diverging pathways, to V2–5, MT, and many other areas with varied functions. The rest go via the superior colliculus in the thalamus to various other cortical and subcortical areas. It has long been known that damage to the eyes, thalamus, and V1 produce blindness, so these early parts are necessary for conscious vision (though see Chapter 8 on blindsight), but may not be sufficient. Conversely, patients with activity in V1 but no connections to higher areas may report no visual experience, and this applies to other senses as well.

Patients in a ‘persistent vegetative state’ (PVS), are described as awake without awareness, somewhere between ‘coma’ (in which a patient has closed eyes and is unresponsive to any stimulation) and a ‘minimally conscious state’ (in which there is some responsiveness or inconsistent signs of consciousness). Belgian neurologist Steven Laureys (2005) tested several patients with what would normally be painful electrical stimulation. This produced activity in the brainstem, thalamus, and primary somatosensory cortex, but not higher up the pain matrix in parietal lobes and anterior cingulate cortex. Similarly, with loud sounds primary auditory cortex is activated, and with flashing light primary visual cortex, but in neither case is there activity in higher association areas (Di et al., 2008). Laureys concluded that PVS is due to disconnection between the primary sensory areas and the fronto-parietal network and that ‘neural activity in primary cortices is necessary but not sufficient for awareness’ (2005, p. 558).

Other evidence comes from studies in which cells in V1 are shown to adapt to invisible stimuli, and from the fact that V1 is suppressed during dreaming sleep even though vivid visual dreams are reported. Studies using single-cell recording in monkeys show that cells in V1 cannot tell the difference between movement caused by eye movements and that caused by movement in the scene, whereas cells higher in the visual hierarchy can – as they must if you are not to think the world has moved every time you move your eyes. From this and other evidence, Koch concludes that ‘While V1 is necessary for normal seeing – as are the eyes – V1 neurons do not contribute to phenomenal experience’ (2004, p. 105).

This might seem a curious conclusion: how can V1 be both necessary for normal seeing and not contribute to phenomenal experience? The underlying assumption here seems to be that most of what goes on in the nervous system is unconscious, and ‘only a fraction of all sensory data pass into awareness’ (Koch, 2004, p. 170). All the unconscious stuff is a necessary precursor of conscious experience but isn’t directly responsible for consciousness. This is clearly a Cartesian materialist description – relying on that magic difference between conscious and non-conscious processes – and it leaves untouched the question of what it might mean for the physical activity of neurons to ‘pass into awareness’.

The default position in NCC research is simply to ignore these questions. For example, a recent meta-analysis of whole-brain fMRI studies contrasting conscious with subliminal visual processing concluded that the NCCs of visual consciousness comprise ‘a subcortical-extrastriate-fronto-parietal network encompassing inferior and middle occipital gyrus; fusiform gyrus; inferior temporal gyrus; caudate nucleus; anterior insula; inferior, middle, and superior frontal gyrus as well

as precentral gyrus; precuneus; intraparietal sulcus; inferior and superior parietal lobules' (Bisenius et al., 2015, p. 180), linking the results to Stanislas Dehaene's global neuronal workspace theory (see Chapter 5). Having earlier supported such a broad fronto-parietal network, Koch's position has switched to a more restricted 'posterior cortical hot zone' for consciousness (Koch et al., 2016). But again, no mention is made of the hard problem or how any of the areas in these impressive lists actually relate to consciousness beyond correlating with reports of seeing the visual stimuli.

Another example concerns eye movements and conscious perception. We may assume that we are usually conscious of whatever our eyes are fixated on, but Miriam Spering and Marisa Carrasco (2015) present converging evidence that eye movements and reportable conscious perception are not tightly linked at all. Rather than being the exception, dissociations between reported experience and eye movements may be the norm. For example, when watching a magician pretending to throw a ball in the air, observers' eye movements do not track the imagined ball, but remain fixated on the magician's face whether or not they experience the illusion. This is presumably because they are watching the magician's direction of gaze (Concept 3.1). The authors suggest that the 'access of the motor system to visual information that does not reach awareness may help manage limited bioenergetic resources' (p. 256). This assumes that consciousness requires extra energy and is an optional addition to functional perception. And this of course raises the questions of why consciousness exists at all, and what difference in the brain (or elsewhere) could make it arise or not.

These examples give a sense of the major methodological and philosophical questions which underlie all research in this area, but which tend to go unaddressed. Reflecting on this state of affairs, Estonian neuroscientists Jaan Aru and Talis Bachmann (2015, p. 1) suggest that despite all the studies conducted over the past twenty-five years, 'it is not clear how much of this research is directly relevant for understanding the neural basis of conscious experience'. Their view is that 'many studies using various experimental paradigms have relied on the contrast between trials with and without conscious perception, but [that] this contrast is not selective for revealing the NCC' (2015, p. 1) because the processing that is studied may in reality either precede or follow from conscious experience rather than directly correlating with it. Along similar lines, Ralph Adolphs, a professor of psychology, neuroscience, and biology at Caltech, points out that while the conceptual problem of consciousness (the hard problem) is notoriously baffling, the methodological problem of consciousness should also not be underrated:

it is very hard to see how the neural correlates of a conscious experience can be separated from everything that accompanies such a conscious experience (our own access to it required for reporting it, antecedent events that make the experience possible, and other events that blur into constitutive components of consciousness).

(Adolphs, 2015, p. 174)

That is, when we think we are studying the correlates of conscious experience, we might actually be studying the correlates of the correlates of consciousness, or the correlates of the precursors of consciousness, or the correlates of access to consciousness . . .

'it is very hard to see how the neural correlates of a conscious experience can be separated from everything that accompanies such a conscious experience'

(Adolphs, 2015, p. 174)

Another criticism is that research has ignored the distinction between phenomenal consciousness and access consciousness. Block's distinction (2007) means that there are two NCCs, not just one: the NCC for phenomenal consciousness and the NCC for access consciousness. The philosopher Benjamin Kozuch (2015) floats the idea that absences in self-report do not necessarily mean that the content represented by the brain area in question is absent from experience. Instead, it might mean that that content is cognitively inaccessible even if someone is conscious of it. Kozuch insists on the possibility that 'one could have a conscious mental state and yet not know it' (2015, p. 146) – that is, have P- without A-consciousness. In his view, people say they are looking for the 'correlates' of consciousness when what they really want to find its 'basis', i.e. that which is minimally sufficient for consciousness, excluding any non-essential correlates.

This way of thinking means accepting two key principles: the P/A-consciousness distinction and the notion that different brain areas 'represent' distinct 'contents'. Alva Noë and Evan Thompson criticise the second of these principles, outlining problems with what they call the 'matching-content doctrine': the belief that 'the first task of the neuroscience of consciousness is to uncover the neural representational systems whose contents systematically match the contents of consciousness' (2004, pp. 3–4). They challenge the majority of neuroscientists for believing that there must be, first, a minimal neural substrate sufficient for making experiences happen, and second, a one-to-one mapping between that and the content of the conscious experience. Chalmers (2000) calls this the *content NCC*, which he says is the 'central case' of an NCC, and Crick and Koch work with this kind of idea when they say that 'Whenever some information is represented in the NCC it is represented in consciousness' (1998, p. 98).

Noë and Thompson (2004) list a number of reasons why the matching-content doctrine doesn't make sense: they argue that perceptual content is structurally coherent, intrinsically experiential, active, and attentional, and exists at a personal, not a sub-personal, level, none of which can be said of neural activity. One of the defining qualities of perceptual experience, for example, is that it is always from a point of view: 'Animals and persons experience the world as laid out before them, but the neurons do not' (p. 16). In other words, neuroscientists are falling for the mereological fallacy. Similarly, as we saw in Chapter 3, occluded parts of an object, like portions of a cat hidden behind railings, seem to be perceptually present even though you can't actually see them. They conclude that neuroscience needs to get away from the ideas of correlation and constitution that define work on the NCCs.

'all qualia are experienced by the brain, and none are reachable objectively from outside that embodied brain'

(Feinberg and Mallatt, 2016, p. 225)

'there can be no match between the content of neural representational systems and the content of experience'

(Noë and Thompson, 2004, p. 88)

COMPETING FOR CONSCIOUSNESS

With or without the help of philosophy, the search for the NCCs has been continued in many forms. The experiments described earlier in the chapter were carried out with monkeys because you cannot ethically insert electrodes into living human brains, but developments in neuroimaging have made it possible to do equivalent studies in humans. Lumer, Friston, and Rees (1998) used fMRI to detect changes during binocular rivalry. Participants wore stereoscopic glasses and were presented with a red drifting grating to one eye, and a green face to the other.




FIGURE 4.5 • (a) Brain areas showing greater fMRI activity during perceptual dominance of the face compared with periods during which the face was unseen are shown as see-through projections onto lateral (left) and horizontal (right) representations of standard stereotactic space (precise 3D positioning). (b) Activity maps during face perception in selected coronal sections, overlaid onto normalised anatomical MRIs. Activity is shown in the fusiform gyri (left), middle and inferior frontal gyri (centre), and dorsolateral prefrontal cortex (right). Distance from the anterior commissure is indicated below each coronal section (Lumer, 2000).

They pressed keys to indicate which they were consciously seeing. These very different stimuli allowed the experimenters to investigate activity in brain areas specialised for the analysis of faces, as well as for colour and form. They looked for areas where activity was significantly higher during periods when the face was seen than when the grating was seen. They found this bilaterally (on both sides of the brain) in occipito-temporal areas of the ventral pathway, including some parts of the fusiform gyrus that are known to be involved in the processing of faces. In addition, they found many prefrontal areas where activity correlated with the image seen.

They also investigated the flipping process itself. To do this, they recorded a participant's series of key presses and then played back the same sequence of images to the participant. They could then compare brain activity for exactly the same sequence of images, but with the important difference that in one case the flipping occurred spontaneously, while in the other it was predetermined. Any differences in activity between the two should then reveal which areas of the brain are involved in causing the flip to occur. Such differences were found in many areas, including parts of the parietal and frontal cortex that had previously been implicated in selective attention (Lumer, 2000). These results add to the impression that conscious visual experiences are correlated not with activity in V1 or other early parts of the sensory pathways but with more central areas.

Other questions we can ask about binocular rivalry are more theoretical. As long ago as 1901 the theory was proposed that neurons encoding the two versions act on each other with reciprocal inhibition. The currently dominant version inhibits the suppressed version, but the cells involved begin to habituate, weakening the inhibition and allowing the suppressed version to win out again until the same process happens in reverse. This would predict that sensitivity in the dominant eye should gradually decrease while that in the suppressed eye increases. Alais, Cass, and O'Shea (2010) investigated this by presenting a brief probe stimulus (a change in contrast at either the bottom or the top of the image) at random times, on average every three seconds, and timed how quickly participants responded to it. The results clearly showed that at the start of a dominance period, probe sensitivity is higher for the dominant image, and lower for the suppressed image, but that this difference reduces towards equal sensitivity by the end of the epoch. This supports the 'adapting reciprocal inhibition model'. It is interesting to note that there are no corresponding changes in the experience of rivalry: although some people seem to briefly see composite images at moments of transition, the dominant stimulus doesn't seem to fade before being replaced by the suppressed one.

Looking at what the suppressed eye does, these experiments also showed that observers are apparently able to respond to a stimulus they are not consciously aware of. One explanation is that the probe presentation itself causes a reversal of dominance that makes the suppressed image detectable, an idea compatible with Dennett's multiple drafts theory of consciousness (see Chapter 5), in which an answer to the question 'what was I conscious of?' is created only by how that experience is probed (by an explicit question, or a particular task, for example). More generally, this finding might indicate that visual awareness manifests either after or at the same time as the planning of motor responses to the probe (Baker, 2010), meaning that it cannot be playing a role in that planning.

Work on binocular rivalry has also been incorporated into a hierarchical model of levels of processing, based on whether the suppressive effects of one temporary blinding method (e.g. binocular rivalry, backward masking, attentional blink) functionally precede or follow those of another method. The neuroscientist Bruno Breitmeyer (2015) locates binocular rivalry at the very lowest, i.e. earliest, level of this hierarchy. He notes that we must bear in mind that this functional hierarchy does not necessarily map readily onto cortical anatomical levels of processing – as a vastly complex network rather than a neatly serial processor, the brain is far too complicated for this to be possible.

Breitmeyer's mapping of the 'functional hierarchies of unconscious processing' can be applied to contexts like the reading process. When reading, we are able to process information not just about the word our eyes are currently fixated on, but also about words we have not yet read. Though we may not be able to identify subsequent words in a sentence before reading them, there may be 'unconscious previewing' of those words with measurable priming effects (Prioli and Kahan, 2015). Breitmeyer (2015) remarks that 'a word stimulus can reveal a type and level of unconscious processing that is higher than that of its basic visual features such as the orientations or curvatures embedded in its graphemic structure' (p. 240). He claims that all visual processing in the subcortical retinal and LGN processing levels is unconscious, but that at the cortical level where processing gets more widely distributed in complex interactions among 'bottom-up', 'same-level', and

'top-down' connections, the story is quite different. This comes close to a dynamical systems theory, to which we will return in the next chapter.

The research on binocular rivalry tells us many interesting things, and ties into major current debates on brain function and structure. But serious problems remain. First, these results provide only correlations, with all the ambiguity we have seen that term to entail. Kanwisher (2001) has suggested that factors including the strength of neural activation and connectivity between neural representations and other parts of the brain may be important in providing the added extra that distinguishes mere necessary conditions from stronger sufficient conditions for consciousness. Some first steps have been taken towards establishing causal connections too, however. For example, Afraz, Kiani, and Esteky (2006) trained monkeys to categorise images as 'face' or 'non-face' and then stimulated clusters of neurons in the ventral stream while they looked at ambiguous images. The monkeys were more likely to indicate 'face' when face areas were activated, suggesting that these areas play a causal role in the act of recognition as well as merely correlating with it. They also found effects of degree of face selectivity and neuronal cluster size in the stimulated sites, and effects of the precise timing of the stimulation. Similar research with humans has found that stimulation of face-selective areas in right fusiform gyrus causes changes in the conscious perception of faces, whereas stimulation of the left fusiform gyrus causes non-face-related visual changes (Rangarajan et al., 2014).

Even when causal links between brain activity and conscious experience are demonstrated, however, they still do not touch the central mystery or help to remove the magic from the 'magic difference'. They do not explain how consciousness could be 'generated' in one place rather than another, how it could 'arise' at one level of processing and not another, or what it means for some processes to be 'qualia-laden' while others are not.

A possibility worth considering is that the whole enterprise is misconceived. For example, if vision is a grand illusion, then there is no 'vivid representation in our brains of the scene directly before us' (Crick, 1994, p. 207). So looking for its neural correlates is doomed to failure (Blackmore, 2002). If we challenge some of the other common metaphors of neuroscience, we might come to the same conclusion: maybe we should call off the search for the correlates of the 'contents of consciousness', because consciousness is not a container.

A different way of interpreting the same data is to imagine that the quality of any experience depends on multiple processes and brain areas. Perhaps a complete integrated system is needed to have the kind of complex, personal, reportable experiences we usually call 'conscious'. This might be in the form of a 'dynamical brain signature' (Lutz et al., 2002) or a certain amount of 'integrated information' in the physical (neural) system (Tononi, 2004; see also Chapter 5). Some, like Andy Clark, Alva Noë, and Francisco Varela, would go further and say that a complex environment or world is needed as well as a whole body and brain. In this case the findings are still fascinating in telling us which brain areas are necessary and/or sufficient for reportable experiences, but they need not indicate that there is a neural location of awareness or that some brain areas 'generate qualia' while others do not.

Maybe we need to make our methods of inquiry even more interdisciplinary to do justice to all these interconnections far beyond the brain, including how they vary from person to person. We will explore this idea in detail in Chapter 17, considering

'there seems to be a magic difference between conscious and unconscious processes'

(Blackmore, 2012)

options like the relatively new field of neurophenomenology, but for now we turn to a different comparison, the differences in brain function between being conscious and unconscious, or being in some of the many possible states in between.

UNCONSCIOUSNESS AND THE BRAIN

Imagine you visit an injured friend in hospital and find her lying passively in bed. Her eyes are open, and she seems at first to be awake but shows no signs of awareness. You try to talk to her, but she does not respond, and you have no idea whether she can hear you or not. Is she still in there somewhere? Does some kind of consciousness remain despite the unresponsive body?

You might worry that your friend has ‘locked-in syndrome’, known in French as ‘*la maladie de l’emmuré vivant*’, or being walled-in alive. This terrifying, though rare, condition happens when parts of the midbrain or brain stem are damaged by accident, disease, or stroke, while higher areas are spared. Usually all muscles are paralysed except for the eyes. So, some patients have learned to communicate using special computer interface technology. A famous example is Jean-Dominique Bauby, whose book *The Diving Bell and the Butterfly* (1997) was dictated one letter at a time by blinking his left eyelid – the only muscle he could move. From such accounts we know that there is a fully conscious, feeling person behind the paralysis. But if your friend were ‘locked in’, she would be unlikely to recover any motor function or to live long.

Monsieur Noirtier was sitting in an armchair which moved on casters, and which he was placed into in the morning, and pulled out of again at night. [...] Sight and hearing were the only senses which, like two solitary sparks, still animated this human substance which was three-quarters of the way to the grave; and of these two, only one could still reveal the inner life that animated the statue; and the look which betrayed this inner life was like one of those distant lights which tell a night-time traveller lost in a desert that a living being still exists in this silence and this darkness.

And in these black eyes of the old Noirtier, crowned by black eyebrows, while all his hair, which he wore long and flowing over his shoulders, was white; in these eyes, as often happens with a bodily organ used to the exclusion of the others, was concentrated all the activity, all the skill, all the strength, all the intelligence, that once had been spread across his body and his mind. Yes, the gesture of his arm, the sound of his voice, the bearing of his body, were lacking, but these powerful eyes replaced them all: he commanded with his eyes, he thanked with his eyes; he was a corpse with two living eyes, and nothing was more frightening, now and then, than this marble face burning from above with anger or glowing with joy.

(Alexandre Dumas, *The Count of Monte Cristo* [Le comte de Monte Cristo], Ch. 58, 1845, our translation)

Alternatively, she might be in a 'persistent vegetative state' (PVS), which we touched on earlier in the chapter. Functional neuroimaging makes possible the investigation of brain states in global disorders of consciousness (Schiff, 2007). These states usually occur with damage to higher parts of the brain but not the brain stem, and patients may progress through to a full recovery or remain in PVS, coma, or a minimally conscious state. In PVS, there is activity in the early sensory areas but not in higher association areas, and no apparent sensory consciousness.

Anaesthesia is a more familiar state to most of us, and research on this provides further insight into NCCs. By varying the dose of anaesthetics and observing the effects with neuroimaging, it is possible to explore the transition from consciousness to unconsciousness, and even the loss of specific functions with deepening anaesthesia. Early experiments with PET scans using the anaesthetics propofol and isoflurane showed a global suppression of cortical functioning with increasing doses, but no evidence of any specific 'consciousness circuits' (Alkire, Haier, and Fallon, 1998).

Subsequent research suggests that for many different anaesthetics the suppression may be caused by blocking at the level of the thalamus, or in thalamocortical and cortico-cortical reverberant loops. This would entail a disconnection something like that seen in PVS (Alkire and Miller, 2005), and has led to the thalamus being described as a possible 'consciousness switch' (Alkire, Hudetz, and Tononi, 2008, p. 877). Alkire and colleagues (2008) argue that the breakdown of cortico-thalamic connectivity prevents the brain from integrating information, and link this to the Integrated Information Theory of consciousness that we will come to in Chapters 5 and 6 – a theory which Tononi and Koch have suggested is the best framework for studying the neural correlates of 'consciousness and other brain functions' (2008, p. 239). On this account, consciousness is not all-or-nothing, but increases and decreases along with the integration of information in the brain. This means that there can be gradual 'shrinking or dimming of the field of consciousness', though at a critical concentration of anaesthetic 'the integrated repertoire of neural states underlying consciousness may collapse nonlinearly' (2008, p. 880).

But not all anaesthetics work in this way. For example, ketamine is a dissociative anaesthetic that is also used as a recreational drug, since at low doses it induces changes in body image, distortions of self, and feelings of dissociation from the surroundings. Ketamine produces an increase rather than suppression in cerebral metabolism, and acts as an antagonist of the *N*-methyl-D-aspartate (NMDA) receptor, blocking the normal excitatory effect of the neurotransmitter glutamate. Other anaesthetics act on other parts of this complex, including nitrous oxide or laughing gas, which is a much smaller molecule but has somewhat similar effects to ketamine. This has led to the suggestion that the normal functioning of the NMDA synapse is necessary for consciousness (Flohr, 2000), which means locating the NCC at a molecular level rather than a higher functional level. However, it may be that despite working in a different way these anaesthetics still affect the thalamus, for example by scrambling rather than blocking signals at the level of thalamocortical interactions (Alkire and Miller, 2005).

Although in principle we should be able to understand consciousness by studying its absence, neither the science nor the logic is straightforward. The use of

'anaesthetics seem to cause unconsciousness when they block the brain's ability to integrate information'

(Alkire, Hudetz, and Tononi, 2008, p. 876)

'Consciousness is not a unitary phenomenon but a catch-all term that includes wakefulness and awareness'

(Shushruth, 2013, p. 1758)

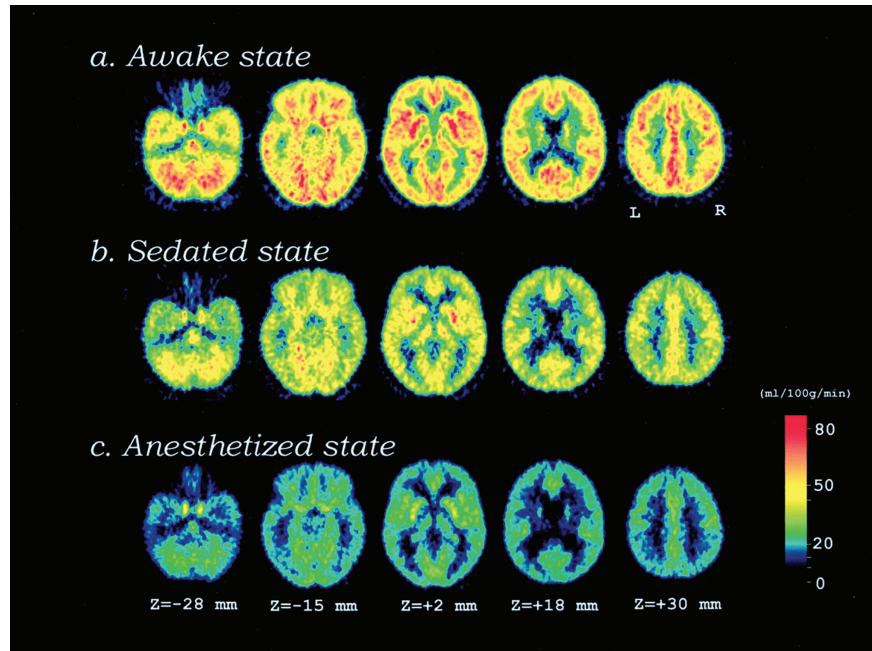


FIGURE 4.6 • A comparison of average PET images of cerebral blood flow in a severely depressed patient for three states: awake, sedated (drowsy with a low dose of propofol), and fully anaesthetised with propofol. Although some researchers thought that anaesthetics might selectively depress specific brain areas associated with consciousness, many studies have shown that the depression of activity is widespread across the brain (Ogawa et al., 2003).

anaesthetics in understanding the neural basis of consciousness raises the important question of what we mean when we say they cause unconsciousness: do we mean they take away wakefulness or awareness? What if the two are not the same thing? (Shushruth, 2013).

Abolishing consciousness is not like pulling out a single component or switching off a light. Nor, as we have begun to see, is the conscious–unconscious distinction necessarily an all-or-nothing difference, in terms of either brain function or experience. Just as we can identify what seems like a spectrum of awareness from coma to minimal conscious state or from heavily to lightly anaesthetised, we can also talk about everyday cognitive processes as involving more or less awareness (Chapter 8). Highly complex perceptual and learning processes go on all the time apparently without awareness. For example, you have no conscious knowledge of how you judge distances, recognise objects from unfamiliar angles, or make aesthetic judgements. You once learned the quickest route to the shops, the name of the highest mountain the world, and the dates of the Second World War, but if we asked you now you would not be able to tell us when you learned them. This is called ‘source amnesia’ and is related to ‘cryptomnesia’ or ‘unconscious plagiarism’, in which people falsely believe they invented an idea themselves when in fact they learnt it from someone else.

Learning a new skill offers a good example of how blurry, and moveable, the boundary may be between conscious and unconscious. You may once have struggled to learn a new language, to ride a bike, or to master the skills of cooking, but now you do these things easily and without paying them much attention.

Skills that were once done with conscious effort have become less conscious, or 'automatised', with time. Reading is a good example. When you first learned to read, every word was difficult and you were probably conscious of each letter, but now you read quickly and with no awareness of individual letters. Bernard Baars (1997a) suggests, as an example of his method of contrastive analysis, that you turn the book upside down and try reading it like that, forcing yourself to go back to a slower and more deliberate kind of reading. So, what does this mean for how the brain changes in the more and less conscious kinds of reading?

Baars correctly predicted that a brain scan would show much more activity in a difficult and more conscious task than in the routine or automated one. A study using fMRI compared a controlled search task with a highly practised and automated task and showed that controlled processing involves a large network of domain-general brain areas (including ACC, preSMA, DPMC, and others), while in automatic processing the control network drops out, leaving activation in only sensory areas (Schneider, 2009).

Dual-process theory is a common model for contrasting 'automatic processing' with the slower and more effortful 'controlled processing' (Kahneman, 2011). There is plenty of evidence for the distinction (see Chapter 8), but the terms may lead us astray. They suggest a kind of unintelligent under-mind that plods along doing boring useful jobs for the truly intelligent conscious mind or striving self. As Wegner points out, the term 'controlled' implies 'a fatal theoretical error – the idea that there is a controller' (2005, p. 19). In his view the controller is an illusion created by mental mechanisms including controlled processes: the controller is an effect, not a cause (p. 20). So, what is unconscious processing? Is it a useful concept, or is the whole idea that we can separate conscious processes from unconscious processes ultimately misguided? As Nancy Kanwisher puts it,

the fact that we can obligate subjects to produce a binary response should not fool us into thinking that their internal state itself is binary or that there is anything important or fixed about the particular threshold the subject uses. Indeed, anyone who has been a subject in a psychophysical experiment will be familiar with the uncomfortable feeling of having to force an unclear and inchoate perceptual experience into one of a small number of discrete response categories.

PROFILE 4.1

Christof Koch (b. 1956)



Known for his multi-coloured clothes and hair, Christof Koch was born in Kansas, but grew up in the Netherlands, Germany, Canada, and Morocco. He originally studied physics and worked at MIT before moving to the California Institute of Technology to run his own K-Lab. He is now also Chief Scientific Officer at the Allen Institute for Brain Science in Seattle, where he aims to create three-dimensional atlases of the mouse and human brains that contain anatomical and genomic data, because 'to understand consciousness, we need to be able to image the activity of millions of individual neurons at the same time'. Koch collaborated with Nobel laureate Francis Crick from the late 1980s until Crick's death in 2004, writing numerous papers and developing a 'framework for consciousness' that guided their search for the neural correlates of consciousness. He first worried about consciousness when he was 18 and in pain: it's just action potentials and ions sloshing about – why should they hurt? Asked how his studies of consciousness have affected his life, Christof said, 'I've stopped eating the flesh of most animals'. He finds ecstasy in running and climbing mountains and once took a solitary mountain hike to convince himself that there really is freedom of action.

(2001, p. 103)

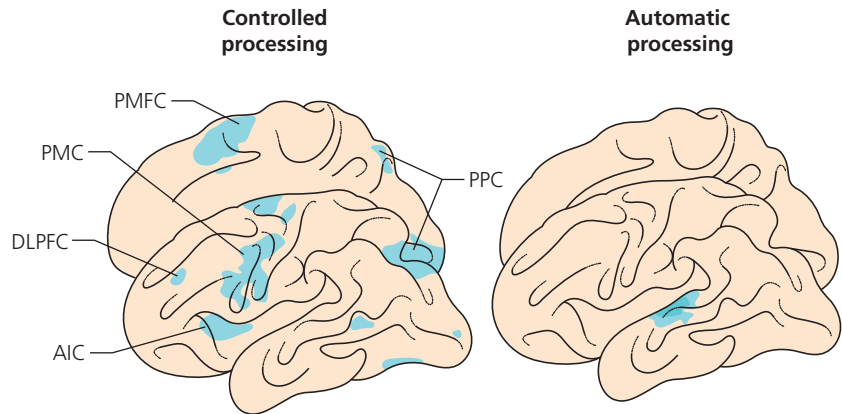


FIGURE 4.7 • (Left) Controlled processing areas activated during activation of cognitive control network regions during visual face and auditory search. Activated areas: ACC/pSMA, DLPFC, IFJ, AIC, dPMC, and PPC. (Right) Automatic processing after extended search for auditory targets. Here the control network has dropped out and only the sensory areas remain active, processing the stimulus via automatic processing (from Schneider, 2009).

Other models of the relation between consciousness and unconsciousness try to avoid the problem of arbitrary oppositions by allowing for both a binary and a graded relationship between the two. One of these is the level-of-processing hypothesis (Windey, Gevers, and Cleeremans, 2013), which suggests that the transition from unconscious to conscious perception is influenced by the level of processing imposed by task requirements.

One last example to consider is the neural basis of something none of us welcomes: a particularly insistent, and embodied, experience – pain.

PAIN

'Pain is always subjective'

(International Association for the Study of Pain, 2011)

Pain hurts. But what does that mean? The all-too-familiar experience of pain raises, in stark form, fundamental questions about NCCs. On the one hand, pain is subjective. The International Association for the Study of Pain defines it as 'an unpleasant sensory and emotional experience associated with actual or potential damage, or described in terms of such damage', and adds, 'Pain is always subjective.'

Perhaps you suspect that your friend who complains at the slightest hint of pain is just a wimp, but how can you know? Just as we cannot know whether your red qualia are just like mine, so we cannot know just how bad someone else's pain really feels. Although genuine facial expressions of pain are – like genuine smiles and laughter – hard to fake, your friend still might either be being terribly brave in the face of agonising pain, or be being pathetic in the face of minor discomfort.

On the other hand, pain correlates with neural events (Chapman and Nakamura, 1999). When someone is injured, numerous chemical changes take place, and signals pass along specialised thin, unmyelinated, neurons called C-fibres to the spinal cord, thence to the brain stem, thalamus, and various parts of the cortex including somatosensory cortex (the precise location depending on where the injury was) and anterior cingulate cortex (ACC). Interestingly, the correlation between the amount of pain experienced and the amount of activity in these

areas turns out to be rather close, with fMRI and PET studies showing larger areas of activation in ACC when pain is rated as more intense.

We all know that pain feels different when it is unexpected rather than self-administered, and worst of all when it's dreaded. This too shows up in ACC. Studies using fMRI have shown that activity in posterior ACC increases with externally applied pain but not with self-administered pain, while activity in perigenual ACC is the reverse (Mohr et al., 2009). All this suggests that there are reliable neural correlates of both the type and amount of pain someone is experiencing.

But what does this correlation mean? Does the neural activity *cause* the subjective experience of hurting? Does the subjective pain *cause* the neural activity? Are both caused by something else? Is pain in fact nothing other than neural activity? Or have we perhaps so misconstrued the situation that we are led to ask impossible questions?

Hold out your bare arm and give it a really good pinch. Now consider this unpleasant feeling. What is it like? While you can still feel it, ask the questions above. Do any of these possibilities really seem right?

'and then of course I've got this terrible pain in all the diodes down my left hand side'

(Marvin the Paranoid Android, in Adams, 1979, p. 81)



PRACTICE 4.1

WHERE IS THIS PAIN?

Look out for any pain you may experience this week, whether a pounding headache or a cut finger. Now look straight into the pain. Experience it as fully as you can. Ask **'Where is this pain?'**

Is the pain located where the headache seems to be? Is the pain inside the cut? Or is it in your head, or in your mind, or where? Are you anxious about the pain, and if so where is the anxiety? Does the pain move when you focus on it? Does it feel as though pain comes into your consciousness and out again? What does this mean?

Odd things can happen when you stare into the face of pain. Make a note of what happens for you.

It might help to explore another question. Where is this pain? Common sense says it is in your arm, which is certainly where it *seems* to be. Identity theorists would locate it in the brain, or perhaps also in all the C-fibres and other activated parts of the nervous system. Dualists would say that it is in the mind and therefore strictly has no location. There are other possibilities too. For example, British psychologist Max Velmans (2009) uses this question to help explain his 'reflexive model of consciousness', in which all experiences result from a reflexive interaction of an observer with an observed (Chapter 17). He rejects both dualism and reductionism and claims that the experienced world and the physical world are the same thing, as looked at from either a first-person or a third-person perspective. On this model the pain really is in your arm.

*Pain - has an Element of Blank -
It cannot recollect
When it begun - Or if there were
A time when it was not -*

(Emily Dickinson, *The Poems of Emily Dickinson*, 1999 [1890], pp. 339–340)

But what if you have no arm? Amputees who experience phantom limbs sometimes suffer excruciating pain in a knee, elbow, or finger that doesn't physically exist. Their pain feels as clearly physically located as yours does.

We began with 'pain hurts', but perhaps we should say 'pain hurts me'. What makes pain painful is the fact that I don't like it; that it's my pain and I wish I didn't have it. Can there then be pains without selves who feel them? And if a self is needed, just how much of a self, and what could the NCs of those necessary selves be?

*'the neural pattern of
you knowing, which is
just another name for
consciousness'*

(Damasio, 1999, p. 73)



PHANTOM PHENOMENA

After losing an arm or leg, more than 90% of people experience a vivid 'phantom limb' that can last for years or even decades. There are also reports of phantom breasts, phantom jaws, and even phantom penises that have phantom erections. Phantom legs can be cramped into uncomfortable positions and hands clenched so hard that the fingers seem to be cutting into the hand. The pain can be excruciating and terribly hard to treat (Melzack, 1992; Ramachandran and Blakeslee, 1998).

The thought of pain in a non-existent limb is so odd that when Silas Weir Mitchell coined the term 'phantom limb' in 1871, after thousands who fought in the American Civil War had had limbs removed from injury or gangrene, he wrote anonymously for fear of ridicule. But it was later studied by Merleau-Ponty as a way of delving into the problems of mind–body-dualism. So where is the pain, and what causes it? An obvious theory is that damaged nerves in the stump send signals to the brain which, wrongly, assumes

CONCEPT 4.2

Consider the case of a dog whose spinal cord has been severed. If a painful stimulus is applied to its leg, the dog shows no signs of distress, but its leg automatically withdraws. Occasionally the same thing happens in humans if they have broken their neck or spine. If prodded in the leg they will deny feeling anything, although their leg pulls back. The isolated spinal cord can even be taught to make responses by training it with stimuli that would be painful for a person with no such injury, but which are not felt at all by the paralysed person.

So, does the spinal cord feel the pain? This is not a daft question. The idea of conscious spinal cords may seem silly, but if you reject this idea then you must also reject the idea that simple animals who have only the equivalent of spinal cords (and no human-like brain) can feel pain. There are also problems with the role of pain in learning. Is the actual feeling of pain, or a pain quale, a necessary component of avoidance learning? If you say 'no', you are led to epiphenomenalism and the possibility of pain-free zombies who learn without *experiencing* the pain. If you say 'yes', then in a simple or damaged organism surely the isolated spinal cord does feel pain, even if it is not like pain in a much more complex whole organism.

Euan Macphail (1998) is among those who deny that other animals can feel pleasure and pain, even though they can learn. Antonio Damasio argues that a self is needed for feeling pain. He argues that neural patterns are not enough – for pain to be painful, and to have the emotional qualities it does, you also have to know that you are feeling it.

Knowing that you have pain requires something else that occurs after the neural patterns that correspond to the substrate of pain – the nociceptive signals – are displayed in the appropriate areas of the brain stem, thalamus, and cerebral cortex and generate an image of pain, a feeling of pain.

(Damasio, 1999, p. 73)

This next stage is also in the brain. It is ‘the neural pattern of you knowing, which is just another name for consciousness’ (p. 73). This means that the necessary neural correlates for pain are to have both the activity in the pain system, and the neural pattern of self – and both are not just correlates, but causes.

Saying that feeling pain depends on knowing you are in pain allows Damasio to distinguish between self-willed actions and automatic reactions, such as removing your hand from the hotplate before *you* even felt the pain – in this case, there is no pain before the action, only afterwards once your knowledge catches up. But Damasio undermines his own argument by saying that even the first pattern alone ‘generate[s] an image of pain, a feeling of pain’ (Damasio, 1999, p. 73). He claims that the feeling of pain depends on knowing one is in pain, but at the same time he still relies on the more traditional assumption that nociceptive signals alone can be sufficient. Notice also that the neural patterns are ‘displayed’, and that the ‘feeling

that the limb is still there. Accordingly, many surgeons have operated on stumps, performed further amputations, cut the sensory nerves, and even operated on the spinal cord, often without stopping the pain.

A completely different approach was taken by Ramachandran (Ramachandran and Blakeslee, 1998). He reasoned that when we clench our fist, feedback from the hand tells us when to stop, but with no hand there is no such feedback and motor signals to clench keep on going, causing the pain. He positioned a mirror in front of a patient so that he could see his normal hand reflected where his phantom should be. When the patient moved his normal hand, he saw what appeared to be the phantom moving, thus providing the necessary feedback. In about half Ramachandran’s cases the phantom seemed to move and the pain eased. In one case, after practising with the mirror, a painful phantom arm that had lasted ten years completely disappeared. Ramachandran claims to have been the first to ‘amputate’ a phantom limb. Other methods have since used sensory and motor retraining, brain stimulation, and virtual reality (Lenggenhager et al., 2014).

The reason why phantoms can be so persistent is that they are part of our body schema. This ‘phantom body’ is the brain’s simulation of our bodily form that uses touch, vision, and other inputs to keep an updated model of our posture, position, and actions, and is essential to coordinate movement. The basic form of the body schema is innate (Melzack, 1989), so that if a limb is lost its ghostly version remains.

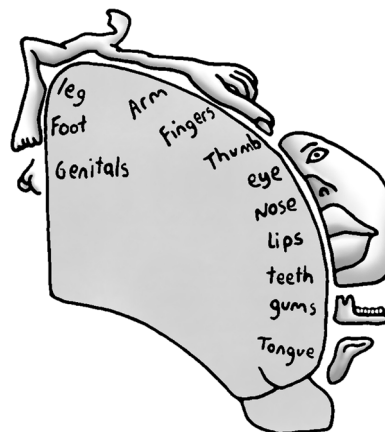


FIGURE 4.8 • The somatosensory homunculus. In the somatosensory cortex each part of the body is represented in a different area. When input from one part is missing, the input from other parts can invade that area. According to Ramachandran, this can explain why amputees sometimes feel real cold on their face as cold in their phantom fingers, or sexual stimulation as a touch on their phantom foot.



FIGURE 4.9 • Ramachandran's mirror box. A mirror divides the open box in half. The patient puts her right hand into the right side of the box and imagines her phantom hand in the left. When she looks into the box she sees two hands. When she tries to move both hands simultaneously a previously frozen and painful phantom is experienced as moving.

WHERE IS THIS PAIN?

'We do not have pains, we get to be pained'

(Humphrey, 2000, p. 13)

'philosophers often ask good questions, but they have no techniques for getting the answers'

(Crick, in Blackmore, 2005, p. 74)

of pain' is equated with an 'image of pain' – notions that imply something watching the displayed image, and hence raise all the problems of the Cartesian theatre. As in global workspace theories, where the contents of consciousness are displayed to the unconscious audience in the rest of the brain, this display is not a magic screening for a psychic homunculus, but is neural activity being made available to other patterns of neural activity. Even so, the problem remains. What is special about this interaction between two neural patterns? What transforms it into a self feeling pain?

The notion of display and the problems of non-identity between brain and mind are avoided by theories that treat sensation as a kind of action. In explaining 'How to solve the mind-body problem', Nicholas Humphrey says that 'sensory awareness is an activity. We do not *have* pains, we *get to be* pained' (2000, p. 13). So, when I feel a pain in my hand, I am not sitting there passively absorbing the sensations coming in; 'I am in fact the active agent' reaching out with an evaluative response, and experiencing this efferent activity. The kind of reaching out characteristic of pain is the movement of pushing away, rejecting, or getting rid of it. In this way he redescribes the 'mind' side of the mystery. 'Thus the phantasm of pain becomes the sensation of pain, the sensation of pain becomes the experience of actively paining, the activity

of paining becomes the activity of reaching out to the body surface in a painful way' (p. 15). The hard problem is, he claims, transformed into a relatively easy problem, although others disagree (see the commentaries following Humphrey, 2000).

Note that Humphrey's theory, although similar, differs from O'Regan and Noë's sensorimotor theory (Chapter 3). They tried to escape from both dualism and the Cartesian theatre by doing away with the idea that perception consists in *representing* the world or the perceiving self. But for Humphrey, the organism 'needs the capacity to form *mental representations* of the sensory stimulation at the surface of its body and how it feels about it' (p. 109). Without this kind of 'inner knowledge', he suggests, sophisticated planning and decision-making simply would not be possible.

We do not know what the necessary and sufficient conditions are for consciousness in general or for particular conscious experiences like experiences of pain. We do, however, know a little about the correlations between brain events and reports of experience. We know, for example, that more activity in the pain system means more intense pain. So, it is natural to wonder – will we one day be able to look into someone's brain and thereby know exactly what they are experiencing? There are hints that this might be possible, but we have also explored some of the reasons why we cannot be confident that the answer to this question will ever be 'yes'.

It seems that even with detailed knowledge of the correlations between brain and experience, we are still far from bridging that gap. In an article called 'Unsolved problems for neuroscience', Adolphs includes the question 'How and why does

conscious experience arise?’ And this he puts into his category of ‘Problems we may never solve’. If we are to have any chance of a solution, we must think comparatively: about cognition across species and across levels of explanation about the brain.

For their part, Noë and Thompson conclude their discussion of the hunt for the NCCs by observing that this quest relies on a specific and controversial notion of conscious content. For them, the moral to be drawn from all this research ‘is that neuroscience, far from having freed itself of philosophy, needs the help of philosophy now more than ever’ (Noë and Thompson, 2004, p. 26). Certainly the claim to be tackling one of the easy problems without begging any conceptual questions seems hard to sustain based on what we have seen so far, and the need for careful research that crosses disciplinary boundaries seems obvious.

In the next chapter we will take another step into the neural labyrinth by asking how the idea of mind maps on to that of brain – or fails to – and what kinds of metaphors may help or hinder our attempts to think about how they fit together.



ACTIVITY 4.1

The rubber hand illusion

This demonstration requires two paint brushes and a dummy hand. The hand can be a life-like rubber model bought specially, as used in the original experiments (Botvinick and Cohen, 1998), or a cheap rubber glove filled with water or blown up and tied like a balloon. This illusion is one of many that provide insight into our body schema (Tsakiris and Haggard, 2005; and see Chapter 15).

The demonstration needs a participant and an experimenter and can be done at home or as a class demonstration. The participant sits and rests their arms on a table, with a screen of some sort to conceal one hand. The dummy hand is then placed in full view, either above or to the side of the real hand. The experimenter takes two paint brushes and gently strokes both the participant’s concealed hand and the dummy hand in exactly the same way at exactly the same time. The experimenter should practise this first and then keep doing it, trying to keep the strokes identical, for a few minutes. The participant, who can see only the dummy hand, should soon begin to feel the sensations as though in the dummy instead of in their own real hand.




FIGURE 4.10 • If the experimenter brushes the hidden real hand and the visible dummy hand in synchrony the dummy hand should start to feel like the participant’s own.

‘neuroscience, far from having freed itself of philosophy, needs the help of philosophy now more than ever’

(Noë and Thompson, 2004, p. 26)

READING



Aru, J., Bachmann, T., Singer, W., and Meloni, L. (2012). Distilling the neural correlates of consciousness. *Neuroscience and Biobehavioral Reviews*, 36, 737–746.

A critique of the ‘contrastive analysis’ method used to study the NCCs.

Crick, F., and Koch, C. (2003). A framework for consciousness. *Nature Neuroscience*, 6, 119–126.

Describes their strategy for studying the NCCs under ten headings relating to their theory of competing cellular assemblies.

Humphrey, N. (2000). How to solve the mind-body problem. *Journal of Consciousness Studies*, 7, 5–112.


Proceeds step-by-step through how to make mind and body line up, proposing an evolutionary theory of sensory awareness as activity.

Kanwisher, N. (2001). Neural events and perceptual awareness. *Cognition*, 79, 89–113.

Describes NCCs based on fMRI, ERPs, and single-cell recordings, discussing which neural events might be necessary and sufficient for perceptual awareness.

Ramachandran, V. S., and Blakeslee, S. (1998). Chasing the phantom. In V. S. Ramachandran and S. Blakeslee, *Phantoms in the brain* (pp. 39–62). London: Fourth Estate.

Fascinating tales of phantom limbs, and Ramachandran’s mirror method for relieving phantom pain.



The theatre of the mind

CHAPTER

FIVE

'The mind is a kind of theatre, where several perceptions successively make their appearance; pass, repass, glide away and mingle in an infinite variety of postures and situations' (1739, I.iv.6). This is how the Scottish empiricist philosopher David Hume described the mind, and the idea of the mind as a theatre has a natural appeal. In Plato's famous allegory of the cave, we humans do not directly see reality but are like prisoners in a dark cave who can watch only the shadows of people outside moving in front of a fire. Two thousand years later, many psychological theories make use of the same metaphor. Yet Hume urged caution: 'The comparison of the theatre must not mislead us', he said. However much we may want to attribute simplicity and identity to the mind, it is a constant flux of transitory impressions: 'They are the successive perceptions only, that constitute the mind; nor have we the most distant notion of the place where these scenes are represented, nor of the material of which it is composed' (1739, I.iv.6). In this chapter we will consider not just those places and materials but the lure, and the dangers, of the theatre metaphor.

I consider that a man's brain originally is like a little empty attic, and you have to stock it with such furniture as you choose. A fool takes in all the lumber of every sort that he comes across, so that the knowledge which might be useful to him gets crowded out, or at best is jumbled up with a lot of other things, so that he has a difficulty in laying his hands upon it.

(Arthur Conan Doyle, Sherlock Holmes in *A Study in Scarlet*, 1887)

INSIDE THE MENTAL THEATRE

What does it feel like being you now? Take a minute to find out.

Although everyone's answer must be slightly different, many people feel they are somewhere inside their head, looking out through their eyes at the world. Indeed, most people adopt a single place inside their body where they feel 'I' am located, and are quite consistent about where this is (Mitson et al., 1976; Limanowski and Hecht, 2011). In one study, participants were asked to explore their sense of self in structured interviews, and 83% confidently located 'the I-that-perceives' in their head, midway between the eyes. This was true for both Chinese and Italians, sighted or blind (Bertossa et al., 2008), and, with different methods, both adults and children (Starmans and Bloom, 2012). In general, research suggests that the most common locations of 'myself' are the upper head or upper torso, with a preference for the head (Alsmith and Longo, 2014) – or, to put it another way, people live either in their brain or in their heart (Limanowski, 2014).



PROFILE 5.1

Daniel C. Dennett (b. 1942)



Dan Dennett, Director of the Center for Cognitive Studies at Tufts University in Massachusetts, studied for his DPhil with Gilbert Ryle at Oxford, and is one of the best-known of contemporary philosophers. Among his many books are *Elbow Room* (1984) and *Freedom Evolves* (2003) about free will; *Darwin's Dangerous Idea* (1996) and *From Bacteria to Bach and Back* (2017) about the evolution of minds; and his challenging *Consciousness Explained* (1991), which demolishes what he calls the Cartesian theatre to replace it with the theory of multiple drafts. He argues for the method of heterophenomenology, rejects zombies as a waste of time, and claims we are all zimboes or conscious robots. He works closely with psychologists and computer engineers, and has long been fascinated by artificial intelligence and robots. He has spent many summers on his farm in Maine, repairing the house, carving wood, making cider, and thinking about consciousness while mowing the hay. Some critics accuse him of explaining consciousness away, but he insists that his really is a theory of consciousness and that, like all good theories, it works like a crane, not a skyhook.

What else do you experience? Does it feel something like this? I can feel my hands on the book and the position of my body, and I can hear the sounds happening around me, which come into my consciousness whenever I attend to them. If I shut my eyes I can imagine things in my mind, as though looking at images hovering somewhere in a mental space in front of or maybe behind my eyes. Thoughts and feelings come into my consciousness and pass away again.

If your experience feels something like this, you may be conjuring up what Dennett (1991) calls the Cartesian Theatre. We seem to imagine that there is some place inside 'my' mind or brain where 'I' am. This place has something like a mental cinema screen or theatrical stage on which images are presented for viewing by my mind's eye. In this special place, everything that we are conscious of at a given moment is present together, and consciousness happens. The ideas, images, and feelings that are in this place are *in consciousness*, and all the rest are unconscious. The show in the Cartesian Theatre is the stream of consciousness, and the audience is me.

Certainly it may feel like this – but, says Daniel Dennett, the Cartesian Theatre and the audience of one inside it do not exist.

Like most scientists and philosophers today, Dennett entirely rejects Cartesian dualism. However, he argues that many who claim to be materialists still implicitly believe in something like a 'centered locus in the brain' where consciousness happens and someone to whom it happens. In other words, there is a kind of dualism still lurking in their view of consciousness. He calls such a belief Cartesian materialism (CM): 'the view you arrive at when you discard Descartes' dualism but fail to discard the imagery of a central (but material) Theatre where "it all comes together"' (1991, p. 107).

It is the view that there is a crucial finish line or boundary somewhere in the brain, marking a place where the order of arrival equals the order of “presentation” in experience because *what happens there* is what you are conscious of.

(p. 107)

Note that the terms ‘Cartesian Theatre’ (CT) and ‘Cartesian materialism’ (CM) are Dennett’s and not Descartes’s. The connection with Descartes is the dualist idea that a specific region of the brain (in Descartes’s case the pineal gland) mediates between the conscious and the unconscious – an idea we might call ‘spatiotemporal pinealism’ (Lloyd, 2000, p. 175). The two terms are also open to various interpretations, and Dennett himself uses them in slightly different ways, which has led to much confusion (Dennett and Kinsbourne, 1992 [including peer commentaries]; Lloyd, 2000). Even so, the central idea is that you believe in the CT if you believe in some kind of literal or metaphorical space or place or stage within which conscious experiences happen, and into which the ‘contents of consciousness’ come and go. You are a Cartesian materialist (rather than being a true materialist, say, or a self-proclaimed dualist) if you also believe that consciousness is not separate from the brain and so there must be some brain basis for this theatre of the mind where “it all comes together” and consciousness happens’ (Dennett, 1991, p. 39).

No one wants to be called a Cartesian materialist; Cartesian materialism has become, as Dennett perhaps intended, a common term of abuse. Nevertheless, says Dennett, this way of thinking is revealed in the way people talk and write about consciousness. People may call themselves materialists and vociferously deny being CMs, but still use phrases that strongly imply the idea of a CT. The CT is a metaphor, but metaphors matter: they let us make sense of abstract things by comparing them with more concrete things (like *my life* with a *journey*), and they are one of our most powerful tools for thought (Dennett, 1991, p. 289; Lakoff and Johnson, 1980/2003). Whichever metaphor we choose, we open up some points of comparison and close down others – often without realising it.

Examples of CM are everywhere, once you start looking: ‘To locate consciousness in the flow of synaptic activity in the brain, we must first locate it in the flow of information processing in the mind’, begins a recent article on the contents of consciousness (Kemmerer, 2015, p. 10). ‘When adopting a descriptive standpoint, even the most cursory examination of the brain reveals a contrast between conscious and unconscious processes’, declare the authors of a paper on the function of consciousness in the nervous system (Morsella et al., 2016, p. 2). More generally, CM is revealed by numerous tell-tale phrases, like saying that a stimulus ‘enters consciousness’, ‘happens outside of consciousness’, or ‘leaps into consciousness’; that some potential ‘content’ ‘comes together in consciousness’, ‘reaches consciousness’ or ‘the level of conscious awareness’, ‘achieves consciousness’, or is ‘unified in awareness’. All these phrases, and many more like them, imply that there is some criterion for what counts as ‘in’ consciousness at any given time,




FIGURE 5.1 • Inside the Cartesian Theatre.

‘Cartesian materialism, the view that nobody espouses but almost everybody tends to think in terms of’

(Dennett, 1991, p. 144)

‘the view you arrive at when you discard Descartes’ dualism but fail to discard the imagery of a central (but material) Theatre where “it all comes together”’

(Dennett, 1991, p. 107)

and that things must be either in or out of consciousness – i.e. on the stage or screen of the metaphorical theatre or not. Even the common phrase ‘the contents of consciousness’ implies that consciousness is a kind of space or container.

If the Cartesian Theatre really does not exist, if ‘consciousness is not a container’ (Blackmore, 2002), then these commonly used phrases must be misleading, and the mistake they depend on may help us understand the confusion surrounding the whole idea of consciousness. If, on the other hand, despite Dennett’s objections, some kind of theatre does exist, we should be able to find out what or where it is. In this chapter, we shall consider the evidence.



PRACTICE 5.1

WHAT IS IT THAT IS CONSCIOUS?

As many times as you can, every day, ask yourself: **‘Am I conscious now? What is conscious?’** Turn your attention away from what you are conscious of to what is conscious of it.

You might be confident that this is ‘me’, but what does that mean? The aim here is to ponder the physical substrate of your experience. Those who object to the mereological fallacy argue that only a whole person can be conscious, not brains or parts of brains, but what does this mean in experience? Try to inquire into your experience now.

Is it your whole physical body that you feel is conscious? Is it your brain? Is it just part of your body or brain? Is consciousness something the brain does? Is it something that could be downloaded into another machine which would still be conscious as you are now? Is consciousness separable from what there is consciousness of?

Are you still conscious now? If so, *what* is still conscious?

THE PLACE WHERE CONSCIOUSNESS HAPPENS

One implication of CM is that there must be a time and a place at which neural processing all comes together to produce a conscious experience – the show in the CT. If this is so, we should be able to find that time and place. We’ll start with what might seem the easier one: the place. So where is it? Let’s take a concrete example of a conscious experience to work with. Right now, please – consciously and deliberately – take a thumb, raise it to your face and press it against the end of your nose. Feel the thumb-on-nose sensations and then let go. It may have felt as though you were sitting in the best seat of your Cartesian Theatre, deciding to do this simple action (or not), and then feeling the sensation. So where did the consciousness happen?

We can easily trace the kinds of neural processing that must have taken place. Reading the instructions would involve activity in much of the visual cortex and in language areas such as Wernicke's area. The oculomotor complex of nuclei would be responsible for moving your eyes as you read, and motor cortex for preparing and executing the skilled action of touching thumb to nose. Frontal areas would be involved in planning and making the decision whether to bother or not. When your thumb touched your nose, parts of the sensory cortex mapped for the hand and face would be activated, and connected with ongoing activity maintaining the body schema (your sense of where your body is in space). In principle, we could examine this activity at any level of detail we wished. But where does the consciousness happen? Two common metaphors imply answers: in one, consciousness is the centre into which experiences come and from which commands go out; in the other, there is a hierarchy of processing with a top where consciousness reigns (see a critical commentary on hierarchies by Feinberg, 2001, pp. 124–125). Global workspace theories of consciousness, to which we will turn later in this chapter, exemplify the first of these; Semir Zeki's hierarchy of micro-consciousnesses, with a single unified macro-consciousness at the top (Chapter 6), is an example of the second. When we talk about neural processing that is 'top-down' (driven by prior goals, expectations, etc.) and 'bottom-up' (stimulus-driven), we might well be drawn to assuming a centre or a top or both.

First, note the obvious point that there is no place in the brain to fit either of these intuitions. As William James poetically put it, there is no 'pontifical' neuron to which *our* consciousness is attached; 'no cell or group of cells in the brain of such anatomical or functional pre-eminence as to appear to be the keystone or centre of gravity of the whole system' (1890, i, pp. 179–180). More than a century later, it is still tempting to think that there must be a centre or a top. But in terms of brain activity, there is no centre and no single top (Zeki, 2001).

To make this clear, we may try to ask which processing is happening on the way in, and which on the way out. Then we might find the middle – where input stops and output begins. This is a reasonable way to think when dealing with a whole organism. After all, light certainly goes into the eyes, and muscles move the arms and legs. So, we can talk unproblematically about input and output. But now we are going right inside the system. Maybe it is just a bad habit of thought, derived from thinking about whole human beings, that leads us to believe that within the brain, too, we can go on looking for 'the middle'. In fact, there can be no middle. Ask yourself whether the activity in Wernicke's language comprehension area is on the way in or the way out, or that in area V1 or V5, or in the temporal lobe. The question makes no sense. There is not a single stream of neural activity coming into a middle and sending a new stream out; there is massive parallel processing. There are feedback loops, complex cell assemblies forming and dissolving, mutual interactions between distant areas, and so on. There is plenty of integration but there is no middle, and therefore also no top (a top being just a particularly privileged kind of middle).

Similarly, there is no special *time* at which consciousness happens. Certainly information comes in first and actions happen later, but between the two there are

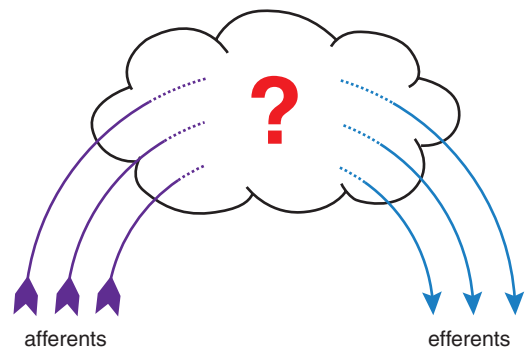


FIGURE 5.2 • Signals come in along afferent nerves, and go out along efferent nerves. So where is the middle, where 'I' receive the impressions and send out the orders? Descartes thought it lay in the pineal gland. According to Dennett, the question betrays a commitment to the Cartesian Theatre. There is no middle, and no 'great mental divide' between input and output (Dennett, 1991, p. 109).

multiple parallel streams of processing, and no magic moment at which input turns into output or consciousness happens, nor any central timing mechanism or 'clock' (Zeki, 2015). In Chapters 6 and 9 we will return in more detail to the timing of consciousness, and the idea that it takes time for consciousness to 'build up'. For now, here is the critical point: we naturally want to ask: 'Which bits of neural processing were the conscious ones and which the unconscious ones? Did I become conscious of my thumb on my nose as soon as it got there?' Dennett argues that even asking such questions betrays a commitment to the CT. They set us off looking for the special time or place where consciousness comes into existence, and that time and place cannot be found.

This argument leads straight back to the hard problem. We assume that in some way all this brain activity is responsible for the powerful feeling I just had that I decided to move my thumb, the thumb did what I told it to, and then I consciously felt the sensation on my nose without being aware of what all those neurons were doing. So, either we have to find an answer to the question 'how does subjective awareness arise from the objective actions of all these neurons and muscle cells?'; or we have to work out what mistake has led us into posing such an impossible question in the first place. Inventing a place and time at which subjectivity happens is not a viable answer.

THE MENTAL SCREEN

In 1971 the American psychologist Roger Shepard published a classic experiment which changed forever how psychologists thought about mental imagery (Shepard and Metzler, 1971). Participants were presented with pairs of diagrams like those shown in Figure 5.3, and were asked to press a button to indicate whether the two were different shapes or different views of the same shape. If you try this, you will probably find that you seem to mentally rotate the objects in your mind's eye. Ask yourself *where* this mental rotation seems to be taking place.

Discussion of such private and unobservable experiences had been banished from psychology by behaviourism, but the importance of this experiment was that Shepard and Metzler made objective measurements. They found that the time taken to reach a decision correlated closely with the time it would actually take to rotate the objects in space. In other words, participants responded more quickly if the object had been rotated only a few degrees, compared with a 180° rotation. Later experiments on imagery showed similar effects. For example, when people are asked to remember a map or drawing and then answer questions such as 'How do you get from the beach to the look-out tower?', the time taken to answer is related to the distance between the starting and finishing points on the map (Kosslyn, Ball, and Reiser, 1978). In other words, it appears that something is happening in the brain that takes time to traverse imagined distance.

The most obvious conclusion is that mental images are like pictures inspected by some kind of mind's-eye function (Kosslyn, 1980), but this was immediately challenged (Pylyshyn, 1973), leading to a long debate between pictorialist and propositionalist (language-like) theories. In essence, the challenge to pictorialism was this: pictorialists correctly observe the similarities between imagery and *vision*, but incorrectly take this to mean that there are *pictures* in the brain, painted on a mental canvas.

'The visual buffer [. . .] is the canvas upon which images are painted; it is the medium that supports depictive representations'

(Kosslyn, Thompson, and Ganis, 2006, p. 18)

Decades later, this great ‘mental imagery debate’ continues (Pylyshyn, 2003; Kosslyn, Thompson and Ganis, 2006), though with some signs of a truce emerging (Pearson and Kosslyn, 2015). When Shepard and Metzler’s experiments were first carried out, no one knew where in the brain the processing was taking place, although there was speculation that the same areas might be used for imagining an object as for seeing it. With the advent of MRI scans and other ways of measuring brain activity, it is now clear that this is correct. When we mentally scan a visual image, similar areas of the visual cortex are activated as when we look at a similar object (Cohen et al., 1996; Pearson and Kosslyn, 2015). But learning more about the similarities between seeing and imagining does not get us very far in adjudicating between pictorialism and propositionalism as competing theories of imagery, since both accept these similarities but draw different conclusions from them.

Both the pictorialist and propositionalist positions have more recently been challenged by a third way of thinking about mental imagery: rather than having either picture-like or language-like images (discrete entities somewhere in the head), in enactivist and sensorimotor theories, we engage in *acts* of imagining. These are closely related to the activity of perceiving the real world (as described in Chapter 3), but enactivist theories hold that in the case of imagining, the sensory exploration is performed without any interaction with the environment. And in the sensorimotor framework, even *potential* for such exploration is enough (Thomas, 2014; Foglia and O’Regan, 2015). These kinds of theory are supported by evidence showing how fundamental action is to both seeing and imagining, for example in a close correspondence between the eye movements we make when we see and when we imagine (Johansson, Holsanova, and Holmqvist, 2006), and even in changes to the thickness of the lens as we see or imagine things close up or far away (Ruggieri and Alfieri, 1992). While pictorialism and propositionalism both rely on the idea that we see and imagine by means of mental representations of the things being seen or imagined, theories within this third camp suggest that action and interaction play much of the role traditionally attributed to representation (Troscianko, 2014, pp. 86–92).

Despite differing interpretations, findings about the timing and the behaviours of imagining and their connections to seeing *do* show that there is something measurable going on when people have private imaginings – imagery is not something mysterious and unamenable to scientific study. They *do not* show either that consciousness is needed to do the imagining, or that there must be a mental screen on which the ‘images’ are projected.

First, mental rotations and other manipulations can happen unconsciously, and indeed do so all the time. When we insert the front door key in the lock, reach out with an accurate grasp to pick up a cup by its handle, or manoeuvre a car into a tight parking space, we deal with rotated imagined objects, but we are

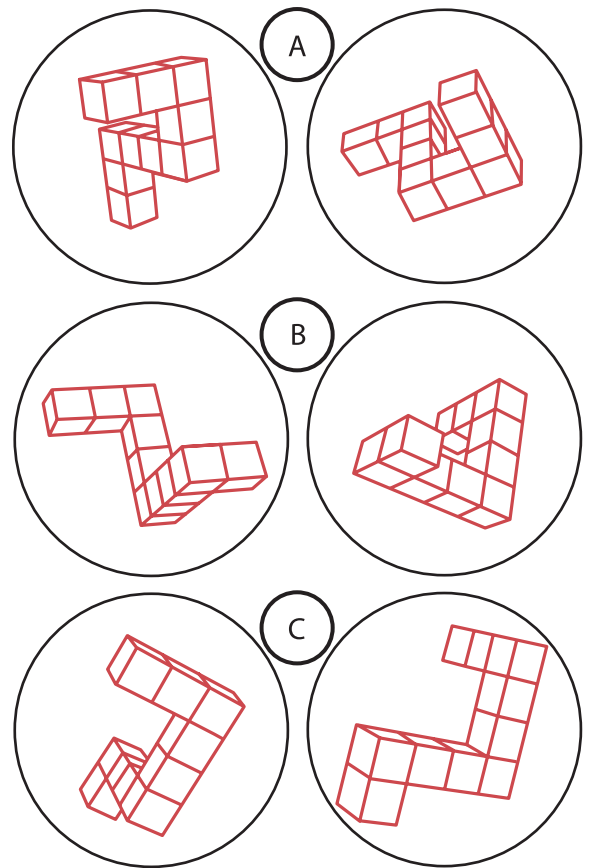


FIGURE 5.3 • In Shepard and Metzler’s (1971) classic experiment, participants had to decide whether the pairs of figures showed the same object rotated or two different objects. The time they took increased with the length of time it would take to rotate real 3D objects.

'mental imagery may involve the same kinds representations [sic] as does vision, and yet in neither case need these representations be pictorial'

(Pylyshyn, 2003, p. 335)

'in the sensorimotor approach, imaging involves being mentally poised to rehearse exploration of an object'

(Foglia and O'Regan, 2015, p. 192)

not necessarily aware of carrying out those rotations. Although imagery is often thought of as quintessentially conscious, similar processes must be going on whether we feel the rotation is done consciously or not.

If you are tempted to think that there must be a mental screen on which the rotated image is projected and that 'you' either do or do not consciously look at the screen and explore its contents, then ask yourself where and what you and the screen could be. If you are a conscious entity looking at the screen, then the classic homunculus problem arises. The inner 'you' must have inner eyes and brain, with another inner screen looked at by another inner you and so on – to an infinite regress.

Crick and Koch claim that there is no infinite regress if the front parts of the brain are 'looking at' the sensory systems at the back. These two areas involve competing coalitions of neurons that interact but not entirely reciprocally, and so give rise to '[t]he illusion of a homunculus in the head looking at the sensory activities of the brain' (Crick and Koch, 2003, p. 124). This would mean that CM really does reflect something about the organisation of the brain. Even so, the nature of this new kind of 'looking' still has to be explained, as does its relationship




FIGURE 5.4 • Imagining pictures in the head means having someone inside who looks at the pictures. That means having someone else inside them looking at their pictures, and another and another, leading to an infinite regress of homunculi.


with consciousness. Crick and Koch hedge their bets by referring to '[t]he (unconscious?) homunculus'. Perhaps all we can safely say about mental rotation is that distributed processing in various areas of the cortex somehow gives rise to the solutions to rotation problems, and either gives rise to or at least correlates with the experience of watching a mental rotation and being able to describe it.

Another example may help. Look around until you find something blue to look at, perhaps a piece of clothing or furniture, or a book or coffee mug. Or close your eyes and imagine a blue cat. Now ask, what is this blueness and where is the blueness located? We know in some detail how colour information is processed in the human brain, and that it must, in some sense, be responsible for the experience of seeing blue. But how?

This is, once again, a version of the hard problem – how does the subjective experience of blueness arise out of all these objective goings-on? An answer that does not work is that the incoming information is turned into a blue picture on a full-colour mental screen for us to look at. There is no single time and place where colour happens. Colour information is distributed through the visual system and used in multiple parallel versions by different brain areas. Even if there were an inner picture, for example in the form of the retinotopic mapping in V1, what would make it blue when all neurons are similar to each other and operate entirely in the dark? As Dennett points out, there couldn't be blue *pigment* in the brain, so could it be *figment*? Of course not. The central mystery is what makes this experience of mine feel so undeniably blue. We cannot solve it by positing a mental screen covered with colour figments and looked at by an inner self. So how can we solve it?

THEATRES THAT ARE NOT CARTESIAN?

The problem that tempts us into imagining a Cartesian theatre is that it seems obvious that we are aware of some of our actions but not others; conscious of some perceptions and not others; have access to some of our desires but not others. So, we have to wonder – what makes the 'magic difference'?



CONCEPT 5.1

SEEING BLUE

- How do we see blue? And why does blue appear the way it does? One problem for consciousness lies in understanding how an experience of seeing blue is related to neural activity in the brain. It may help to think about how colour processing works.

There are three types of receptor in the retina (somewhat misleadingly called red, green, and blue cones) that respond differentially to different wavelengths of light hitting them. Output from the red and green receptors is subtracted to produce one dimension of colour and summed to produce a luminance signal (which then contributes to other kinds of visual processing). This sum is subtracted from the output from the blue cones to produce a second colour dimension. These two-colour opponent processing signals are sent (as rates of neural firing) via the optic nerve to the thalamus and then to the visual cortex. In the visual cortex, some areas use only luminance information and construct edges, movement, and other visual features, while some also use the colour information and incorporate it into processing visual scenes and perceived objects. Output from this processing is then used in further brain areas dealing with associations, memory, and the coordination of behaviours. So, when you look at a blue mug, neurons throughout the visual system are firing at a different rate or in different patterns from how they would fire if the mug were orange.

But where does the *experience* of blue happen? Where are the qualia? Where or when, in all this processing, does the *conscious experience* occur? Theories of consciousness must either:

- 1 Answer the question, for example by proposing a brain area, a special kind of processing, or a feature of functional organisation that is responsible for consciousness.

Or

- 2 Explain why there is no answer.



FIGURE 5.5 • Experienced drivers may find that they arrive at their destination with no memory of having driven there at all. Were they really conscious all along but then forgot the experience (an Orwellian interpretation)? Was consciousness prevented by something else, such as paying attention to a conversation or music (a Stalinesque interpretation)? Were there two parallel consciousnesses, one driving and one talking? Or does the unconscious driving problem reveal the futility of asking questions about what is ‘in’ consciousness at any time? Theories of consciousness account for this phenomenon in many different ways.

In other words, what makes some events *conscious* and others *unconscious*; some *in consciousness* and some *outside of consciousness*?

Let’s take the most familiar example – the unconscious driving phenomenon. You drive on a well-known route, say to work or a friend’s house. On one occasion you are acutely aware of all the passing trees, people, shops, and traffic signals. Another day you are so engrossed in worrying about consciousness that you are completely unaware of the scenery and your own actions. You realise only on arriving at your destination that you have driven all that way unaware of what you were doing. You have no recollection at all of having passed through all those places and made all those decisions. Yet you must, in some sense, have noticed the traffic signals because you did not drive through a red light, run over the old lady on the crossing, or stray on to the wrong side of the road. You applied the brakes when necessary, maintained a reasonable driving distance from the car in front, and found your usual route. So, considering the red light, what makes the difference between its being *in consciousness* and *out of consciousness*?

This is where the Cartesian theatre comes in. We can easily imagine that during each of these journeys the things I was conscious of were on the stage, and all the others were not: the things ‘I’ was aware of were those presented to my mind’s eye, visible on my mental screen at the time, available to ‘me’ to look at, consider, or act upon. But there is no literal place inside the brain that constitutes this theatre. So, what are the alternatives? Some theories keep the theatre metaphor while trying to avoid the impossibilities of a Cartesian theatre. Some throw out all theatre

imagery and try to answer the question another way. Others are more radical and even throw out the idea that things are unequivocally *in* or *out* of the stream of conscious experience, or that there is any such stream at all.

In the rest of the chapter we will give examples of each type, but almost any theory of consciousness can be categorised in terms of its answers (or non-answers) to the big three central questions. 1) Does the theory try to solve the problem of why there is subjective experience at all? 2) Does it try to explain what makes some events *conscious* and others *unconscious*? 3) Does it try solve one or both of these problems by positing a mental or neural theatre? If so, is this a Cartesian theatre? Bear these questions in mind as you assess the ideas that follow.

We'll start with the most explicitly theatrical of current theories. Bernard Baars's global workspace theory (GWT) was first developed in the 1980s, and has since been extended in biocomputational directions (Baars and Franklin, 2009) and into a fully-fledged neural theory (by Dehaene and his collaborators; see below).

Baars begins by pointing out the dramatic contrast between the very few items that are available in consciousness at any one time and the vast number of unconscious neural processes going on. The best way to think about this, he argues, is in terms of a theatre. Focal consciousness acts as a 'bright spot' on the stage, which is directed to different actors by the spotlight of attention, possibly surrounded by a fringe of events that are only vaguely or potentially conscious (Mangan, 2001; Shanahan, 2006). Meanwhile, 'The rest of the theater is dark and unconscious' (Baars, 2005a, p. 47). The unconscious audience sitting in the dark receives information broadcast from the bright spot, while behind the scenes there are numerous unconscious contextual systems that shape the events happening in the bright spot.

GWT is explicitly based on the 'Theater Hypothesis' (Baars, 1988), or 'A theater metaphor and brain hypotheses' (2005a). These entail a 'theater architecture' that includes a 'theater spotlight'. Conscious events happen 'in the theatre of consciousness' or on 'the screen of consciousness' (Baars, 1988, p. 31), and 'Conscious contents correspond to the bright spot on the stage of working memory' (2005a, p. 47).

What makes this theory more than just a loose metaphor, claims Baars, is its grounding in psychology and neuroscience. Backstage are the processes creating the current context, while the bright spot of attention corresponds to the contents of consciousness, and the rest of the stage to immediate working memory



PROFILE 5.2

Bernard Baars (b. 1946)



Born in Amsterdam, the Netherlands, Bernard Baars moved to California with his family in 1958. He trained as a language psychologist before moving into consciousness studies. He says that living with cats makes it seem obvi-

ous they are conscious, with ethical implications for dealing with animals, babies, fetuses, and each other. His well-known global workspace theory was inspired by artificial intelligence architectures in which expert systems communicate through a common blackboard or global workspace. He describes conscious events as happening 'in the theatre of consciousness', where they appear in the bright spotlight of attention and are broadcast to the rest of the nervous system. He advocates investigating consciousness through the method of 'contrastive analysis': comparing closely matched conscious and unconscious events. He was a Senior Fellow in Theoretical Neurobiology at the Neurosciences Institute in San Diego, and co-founded the Association for the Scientific Study of Consciousness, as well as the journal *Consciousness and Cognition* and the online resource *Science and Consciousness Review*. As far as consciousness is concerned, he thinks we are at last beginning to see the light.

• SECTION TWO : THE BRAIN

(Baddeley, 2000). The interactions between these three are based on the idea of a global workspace architecture, first developed by cognitive modelling and common in computational approaches to human cognition. On this view, the brain is structured so that just a few items at a time are dealt with in the global workspace – similar to the 7 ± 2 items conventionally held in working memory. The theatre has numerous inputs from the senses and from the overall context, and connections to the resources of a ‘vast unconscious mind’ (Baars, 1997a, p. 304), such as language, memory systems, and learned skills. According to Baars, all this provides a real ‘working theatre’, with consciousness acting as a gateway, providing global access to any part of the nervous system.

In this scheme, consciousness has very definite effects and functions. It provides access to the mental lexicon, and to autobiographical memory and the self-system. It recruits processors for ongoing tasks, facilitates executive decisions, and enables voluntary control over automatic action routines. According to Baars, consciousness is not an epiphenomenon, nor is it mysterious. It is a working part of the cognitive system. Baars understands the brain as a whole as consisting of

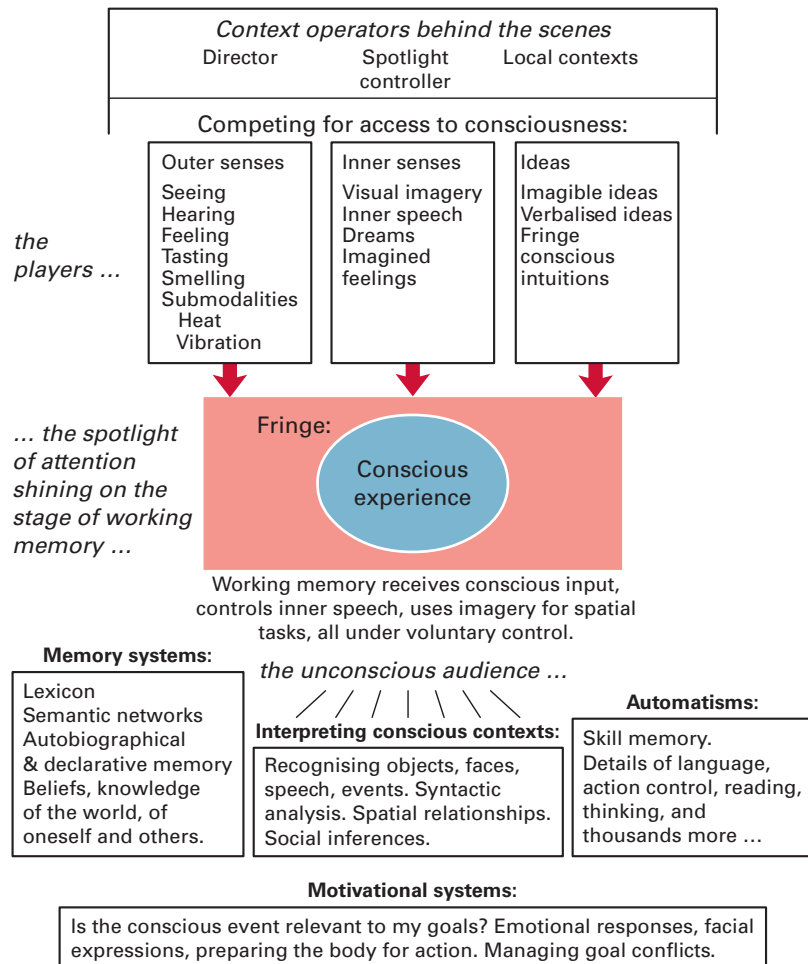


FIGURE 5.6 • Baars's theatre metaphor for conscious experience (Baars, 1997a, p. 300).

multiple decentralised networks (2005a, pp. 47–48), but in GWT consciousness is required to integrate and coordinate these otherwise autonomous networks. So, the brain turns out to be highly centralised after all. How exactly this kind of integrative function might be achieved by virtue of things in the spotlight being conscious is a question not addressed by Baars, and one we will return to in Chapter 11, on the function of consciousness.

On this theory, what makes an event conscious is that it is being processed within the global workspace and is made available to the rest of the (unconscious) system. So, when you drive with full attention, information pertaining to the red traffic light is processed in the global workspace. When your workspace is filled with philosophical speculations and imagined conversations, the red light is no longer in the spotlight on the stage, and is relegated to the fringe or even to the darkness.

Baars's preferred method of investigation is to treat consciousness as a continuous variable, contrasting 'more conscious' with 'less conscious' events, while holding the contents of experience constant. He calls this way of looking for the NCCs 'contrastive analysis', and says it could be adapted to test predictions specifically of GW theories. Experiments using scanning or other methods might be designed to find out what processes in the brain are involved when the same thing is right in the spotlight of focal consciousness, in the less conscious fringe, or outside consciousness altogether.

An example is the fading of words into short-term memory. The same words might be at one time in conscious inner speech (in the bright spot on the stage), then fade into unconscious but easily accessible working memory (still on the stage but out of the spotlight), and then become conscious again (move back into the light) when retrieved, or alternatively fail to be retrieved (leave the stage altogether). Any complete theory of consciousness has to explain the difference, says Baars. Rather than worrying about the hard problem, we should get on with the task of finding out what makes events more or less conscious.

We can now assess GWT by its answers to our three questions: 1) Why is there subjective experience at all? 2) What makes some events *conscious* and others *unconscious*? 3) Does it posit a mental or neural theatre, and if so, is it a pernicious Cartesian theatre?

On the last point, GWT obviously involves theatres, but Baars (1997a, p. 292) argues that 'Working theatres are not just "Cartesian" daydreams' and that fear of the Cartesian theatre is misplaced, for no one believes in a single point at which everything comes together, and his theory does not require it. Yet he does argue for something like a convergence zone somewhere in the brain. He claims that 'there is indeed a place in the visual system where "it all comes together"', and that this may be involved in constructing the global workspace. He likens the visual system to a (rather complicated) staircase, at the top of which 'The brain regions for object recognition appear to be where the contents of consciousness emerge' (in Blackmore, 2005, pp. 16, 13). He adds that the spotlight might correspond to some kind of attention-directing mechanism, and that research on the self-systems which construct inner speech and provide a running narrative on our lives could usefully be guided by the metaphor of a theatre. Despite his rejection of 'Cartesian daydreams', Baars does assume that at any given time some things

'all of our unified models of mental functioning today are theatre metaphors; it is essentially all we have'

(Baars, 1997a, p. 301; also 1997b, p. 7)

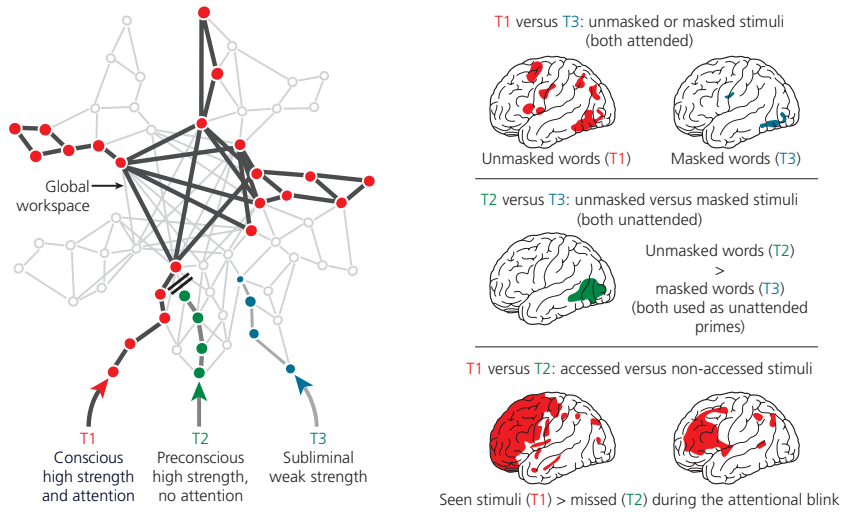


FIGURE 5.7 • (Left) Dehaene and colleagues' scheme for showing two ways in which stimuli can fail to gain access to the global neuronal workspace. (Right) Reinterpretation of neuroimaging experiments in this framework (after Dehaene et al., 2006).

are *in* consciousness, while others are not – the assumption which is, according to Dennett, at the heart of Cartesian materialism.

On the second point, Baars clearly distinguishes between conscious and unconscious events, the difference being whether they are in or out of the GW.

To tackle the even trickier problem of subjectivity it may be useful to consider a related theory, the 'neuronal global workspace' model proposed by French neuroscientist Stanislas Dehaene and his colleagues (Dehaene et al., 2006; Dehaene, 2014). The immediate advantage of this version of the GW idea is that no process or information has to 'become conscious' or 'enter consciousness'. In this version, a collection of specialised unconscious processors compete for access to the limited-capacity global workspace, which is probably dependent on long-range circuits involving the prefrontal cortex, anterior cingulate, and connecting areas (Dehaene and Naccache, 2001). Information can then be widely broadcast to other brain areas, and 'This brain-scale broadcasting creates a global availability that results in the possibility of verbal or non-verbal report and is experienced as a conscious state' (Dehaene, 2009, p. 468).

The crunch here comes with the word 'and', which can be interpreted in two completely different ways. One implies an *and then*: when information gets into the GW, something else very special happens – *and then* it 'becomes conscious' or 'enters consciousness'. If this is the interpretation you favour, then the transformation is magic – or at least completely unexplained. Neuronal GWT cannot account for subjectivity or deal with the hard problem.

The other interpretation is that the 'and' equates the two. That is, being accessible to report simply *is* what we mean by being conscious: subjectivity and access are the same, and there is no hard problem to solve. This is the interpretation that Dennett urges upon us; for him, the hardest part of understanding the neuronal GW thesis correctly is acknowledging that global availability does not cause some

'brain-scale broadcasting creates a global availability that results in the possibility of verbal or non-verbal report and is experienced as a conscious state'

(Dehaene, 2009, p. 468)

further effect, 'igniting the glow of conscious qualia, gaining entrance to the Cartesian Theater, or something like that' (Dennett, 2001b, p. 223). 'Those who harbour this hunch are surrendering just when victory is at hand,' says Dennett (*ibid.*), because global availability just *is* a conscious state. Consciousness is like 'fame in the brain' or 'cerebral celebrity'; fame is not something *in addition* to being well known, nor is consciousness. When Dehaene says in his later book *Consciousness and the Brain* simply that 'consciousness *is* global information broadcasting within the cortex' (2014, p. 13, our italics), he seems to endorse the 'there's nothing extra' idea. But when he continues the sentence – 'it [consciousness] *arises from* a neuronal network whose *raison d'être* is the massive sharing of pertinent information throughout the brain' (p. 13, our italics) – space is opened up again for the other interpretation: that once it's been shared, something else happens to the information to make it conscious.

'the theatre metaphor seems to have outlived its usefulness'

So how should we interpret GWTs? Especially since being given a new neural spin by Dehaene, they have remained popular, but like many if not most current theories of consciousness, they also evade or are ambiguous about the critical question of what actually makes something conscious. This is true of the many spin-outs and implementations of the theory (e.g. Maia and Cleeremans, 2005; Gaillard et al., 2009; Raffone and Pantani, 2010) as well as its main expositions. Thus the British neuroscientist David Rose concludes that 'It is difficult to keep in sight what consciousness actually is in global workspace theories, let alone its source' (2006, p. 222). In particular, GWTs can often raise rather than answer the questions 'is consciousness the cause or the result of access to the global workspace?' and 'is global availability a consequence of or an explanation for consciousness?' (p. 223).

(Rose, 2006, p. 223)

One question to ask is whether the brain is actually organised with a GW, and what precisely this would mean (Dehaene, 2014). If it is, then we must ask whether entry to the GW is a cause or a consequence of something becoming conscious. Either way, we must then decide whether just being in the GW bright spot accounts for the blueness of blue or the feeling of observing your mental images, or whether something more is needed to turn the contents of the GW into subjective experiences. For Dehaene (2014, p. 262), there is no question that 'Once our intuition is educated by cognitive neuroscience', the hard problem will be revealed as non-existent. On the basis of what we have learned about the theatres of the mind so far, do you agree?

WHAT IS CONSCIOUS NOW?

THEORIES WITHOUT THEATRES?

Getting rid of the idea that we need a theatre of some kind – that is, a distinction between things that are in or out of consciousness, or between bits of brain where consciousness does or doesn't happen – is vastly difficult. Some of the simplest ways of doing it are identity theory, which equates conscious experiences with brain activity, and the eliminative materialism espoused by Paul and Pat Churchland (e.g. P. M. Churchland, 1981; P. S. Churchland, 2002), which thinks there is nothing to be explained beyond the material. Paul Churchland is happy to talk about qualia such as the redness of the red light, and even to revel in them as what makes life worth living, but denies there is any special problem of subjectivity. He likes to take lessons from the history of science. 'Electromagnetic waves don't *cause* light; they're not *correlated* with light; they *are* light. That's what light

• SECTION TWO : THE BRAIN

is' (in Blackmore, 2005, p. 54). Similarly with sound and heat. Even though it seems difficult for us now, he – like David Papineau, whom we heard from in Chapter 2 – thinks that in time we'll come to accept that to have a sensation of red

is to have all of your three kinds of opponent processing cells showing a certain pattern of relative stimulation. [. . .] The pattern of activation for red will be, say, 50%, 90%, 50%, across the three kinds of cells.

(2005, p. 55)

There is no need for any theatre imagery here, and the problem of subjectivity is dealt with by claiming identity between neural processes and subjective experiences. There remains a problem with our driving example, though, because presumably the driver attending to the road and the distracted driver who still stopped at the red light would both have had the right proportions of opponent processing cells firing in their visual systems. The difference in their experiences would have to be accounted for in some other way, perhaps by differences in recall when they got to their journey's end. More generally, this account does not explain how we can overcome our intuitions, or advance the science, enough to see how experience *is* brain activity.

'You're nothing but a pack of neurons'

(Crick, 1994, p. 3)

Many other theories of consciousness avoid the controversial imagery of stage and theatre, at least on the surface. These include the most explicitly reductionist theories, like Crick's 'astonishing hypothesis': 'that "you", your joys and your sorrows, your memories and your ambitions, your sense of personal identity and free will, are in fact no more than the behaviour of a vast assembly of nerve cells and their associated molecules' (Crick, 1994, p. 3). The theory involves no explicit theatre imagery, yet Crick compares thalamic control of attention with a spotlight, giving a hint of the theatrical. He claims that brain activities 'reach consciousness', and speaks of 'the seat of visual awareness' (p. 171), of 'the location of awareness in the brain' (p. 174), and of locating the 'awareness neurons' (pp. 213, 224). So arguably, Crick's theory is still a form of Cartesian materialism.

As far as the red traffic light is concerned, Crick's early theory required the right oscillations to bind the features of the red light. His later theory with Christof Koch involves the activation of thalamocortical loops. In both cases (see Chapter 6), the theory requires specific brain processes that correlate with the light being consciously perceived or not.

Other theories avoid the theatre by focusing on massive cross-brain integration. For the Dutch neuroscientist Cyriel Pennartz, for example, consciousness is the solution to 'the brain's representational problem': how to integrate multiple pieces of sensory information 'into a coherent whole that can be immediately recognized, rapidly understood, and acted upon' (2015, p. 10). Pennartz splits up the requirements for consciousness into 'hard' (non-optional) and 'soft' (optional though common). The 'hard' prerequisite of consciousness is an ability to interpret multiple sensory inputs as having particular qualities, meaning, or content – in our example, interpreting all the visual qualities of the red light in tandem with surrounding stimuli, and attributing the meaning 'stop' to them. The 'soft' requirements include projection of interpreted sensory inputs into an external, perspectival space (as

in vision) or body map (as in somatosensation), and the construction of an illusion of 'unity' in consciousness and self-awareness. Pennartz goes on to identify neural candidates for carrying out these functions, including mechanisms for coordinating, binding, and stabilising, and for varying the phase and rate of cell firing.

For Pennartz, the brain has to be understood as operating in a high-dimensional space, with each sensory modality or submodality constituting an additional dimension. But he does not quite manage to explain why the 'representational power' of the brain, vast as it may be, should in itself have or yield a subjective quality. He has to resort in the end to a distinction between 'brain systems for conscious *versus* unconscious representations' (2015, p. 113, original italics) and between the 'neural relationships [that] give rise to' one or other kind of representation (p. 288) – which means that even if the theatre no longer has a stage and a spotlight, it still exists in a Cartesian materialist sense.

As Pennartz points out (2009, p. 733), his view of the importance of multidimensional integration bears some resemblance to the cross-brain broadcasting of information in GWT. In this sense, it also resembles what is currently probably the most popular of all theories of consciousness, integrated information theory (IIT).

IIT was originally proposed by Giulio Tononi in 2004, building on his work with Gerald Edelman (Tononi and Edelman, 1998), and has since been updated several times (Tononi, 2015). The basic principle is that the more 'integrated information' there is in a physical system, the more conscious that system is, and the amount of integrated information is measured by a mathematical variable, Φ (phi).

As with Baars's theory, consciousness is a continuous variable: you can have different amounts of it. In this case, the system becomes conscious (and has free will) if it has a large value of Φ , and a system is more conscious the higher its Φ value. The fact of having a large value of Φ can also help explain the specific qualities of a given conscious experience compared to all the other possible ones: because 'generating a large amount of integrated information entails having a highly structured set of mechanisms that allow us to make many nested discriminations (choices) as a single entity' (Tononi, 2008, p. 224), we experience the red light not simply as the opposite of no light, or of green, but as different from any other possible experience we might have.

For IIT, 'consciousness is integrated information', and 'its quality is given by the informational relationships generated by a complex of elements' (Tononi, 2008, p. 217). This means that IIT need not posit a Cartesian theatre, because any part of the nervous system can in theory contain integrated information. What happens to any informational relationships that exist outside the integrated system? Tononi gives the examples of sensory afferents or cortico-cortical loops implementing informationally insulated subroutines. These 'do not make it into the quale, and therefore do not contribute either to the quantity or to the quality of consciousness' (p. 229). This begs the now-familiar question of why it is that integrated information should 'make it into the quale' (i.e. become conscious experience) and other kinds not – or what a quale actually is, and how you get inside one. So, we do have a theatre of sorts, even though it is the opposite of a spatially localised one. We will come back to other implications and criticisms of IIT in Chapter 6.

Another attempt to bridge the explanatory gap without the help of a theatre appeals to quantum-level processes – that is, processes involving the smallest possible amount of a given physical entity (like a photon or an electron). For the British physicist and mathematician Roger Penrose, to solve the hard problem we need to understand the problem of incompatible explanatory levels. There are two levels of explanation in physics: the familiar classical level used to describe large-scale objects, and the quantum level used to describe very small things and governed by the Schrödinger equation. Both these levels are completely deterministic and computable. The trouble starts when you move from one to the other. At the quantum level superposed states are possible – that is, two possibilities can exist at the same time – but at the classical level either one or other (seeing the light as red or green) must be the case. When we make an observation (working at the classical level), the superposed states have to collapse into one or other possibility, a process known as the collapse of the wave function.

A variety of theories have been developed using quantum physics to try to solve the hard problem (Tuszynski, 2006). Some physicists, notably Eugene Wigner and Henry Stapp, have claimed that consciousness causes the collapse of the wave function. On Stapp's theory, the quantum brain is understood as a 'collection of classically conceived alternative possible states of the brain' which 'all exist together as "parallel" parts of a potentiality for future additions to a stream of consciousness' (2011, pp. 51–52). In this context, nondeterministic consciousness controls deterministic brain activation by an attentional process of choosing between alternatives (Stapp, 2007). This quantum interactive dualism involves a widespread effect and is different from Popper and Eccles' dualist interactionism (Chapter 6), in which the mind intervenes at certain synapses in what would otherwise be a causally complete physical system.

Far from insisting on physicalist explanations, Stapp believes that 'contemporary physical theory demands certain interventions into the physical! The associated causal gap in a purely physically determined causation provides a natural opening to an interactive but non-Cartesian dualism' (2011, p. 116). As a result, he claims that this kind of quantum approach can solve the binding problem and explain the unity of consciousness and the power of free will.

Note that these ideas from quantum physics have inspired many popular and spiritual theories. These include nuclear physicist Amit Goswami's theory of creative evolution in which consciousness is the ground of all being, and spiritual teacher Deepak Chopra's notion of consciousness as fundamental, as a field phenomenon that precedes the quantum field as the origin of the universe (Kafatos et al., 2011). Among others are theories of quantum consciousness, quantum awakening, and the quantum soul (Zohar and Marshall, 2002).

But this is not what Penrose means. Penrose argues that all conventional interpretations of the collapse of the wave function are only approximations, and instead proposes his own theory of 'Orchestrated Objective Reduction', or Orch OR. This new process is gravitational but nonlocal in nature, and hence can link things in widely separated areas, making large-scale 'quantum coherence' possible. This can only happen when the system is isolated from nonorchestrated perturbations in the rest of the environment, so that objective reduction and the hidden noncomputational action it makes possible can be made use of by the system in

a controlled way. This kind of stable isolation is normally possible only at extremely low temperatures.

Where in the brain could such a process, requiring such particular conditions of stability, be going on? Penrose builds on the suggestion first made by American anaesthetologist Stuart Hameroff that consciousness emerges from quantum coherence in the microtubules. Microtubules are, as their name suggests, tiny tube-like proteins found in almost all cells of the body. They are involved in supporting the cell's structure, in cell division, and in transporting organelles within the cell. Hameroff and Penrose propose them as the site of non-algorithmic quantum computing because of their shape and the spiral structure of their walls, and because any quantum-coherent effect within them could be kept reasonably isolated from the outside.

Why is this relevant to consciousness? Hameroff argues that the real problems for understanding consciousness include the unitary sense of self, free will, and the effects of anaesthesia, as well as non-algorithmic, intuitive processing. All these, he claims, can be explained by quantum coherence in the microtubules. Nonlocality can bring about the unity of consciousness, quantum indeterminacy accounts for free will, and non-algorithmic processing, or quantum computing, is done by quantum superposition. So, it is in the microtubules that not only your experience of the traffic light on red is generated, but also your sense that it is 'you' who experiences the red, and you who freely decide to stop the car when you see it.

In 2014, Hameroff and Penrose published new evidence in support of the theory, including a claim to have identified microtubule 'quantum channels' in which anaesthetics erase consciousness, as well as warm quantum vibrations in brain microtubules, and links between the 'beat frequencies' of microtubules and the gamma synchrony found in EEG. Among many commentators on the 2014 paper, Deepak Chopra contrasts mind-first theories that have trouble explaining matter against matter-first theories (including this one) that struggle to explain mind. Hameroff and Penrose respond that quantum events indicate an 'invisible agency (consciousness)' (p. 96) producing intelligent activity at the interface where space-time emerges. Responding to others, they agree that the current model of neuronal functions is inadequate, that parts of the human mind are non-computable, and that without an account of consciousness, the universe remains 'as mysterious as ever' (p. 98).

A sticking point has always been whether quantum coherence could survive in a warm, wet brain. Hameroff and Penrose argue that biology can use thermal energy to drive coherence, while physicist Matthew Fisher (2015) has proposed that the nuclear spins of phosphorus atoms in the brain could allow it to function as a quantum computer. But if the brain is a quantum computer, does this tell us anything about consciousness?

More generally, we might question whether this quantum theory doesn't just replace one mystery (subjective experience) with another (quantum coherence

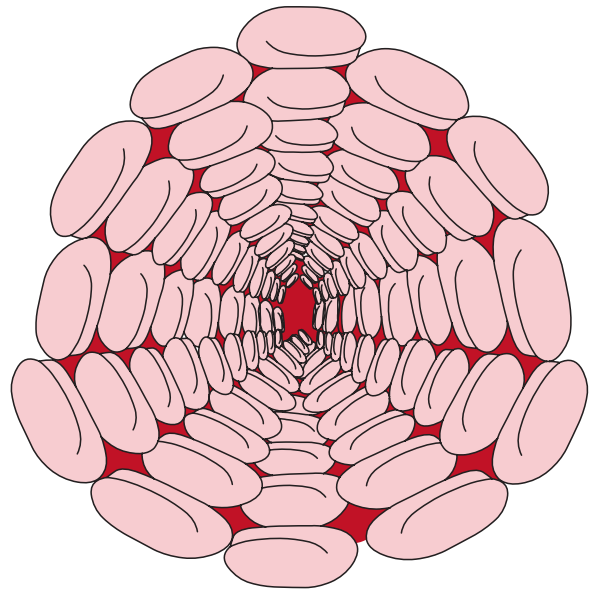


FIGURE 5.8 • Penrose and Hameroff argue that consciousness emerges from quantum coherence in the microtubules. Microtubules are structural proteins in cell walls. They are shaped like a hollow tube with a spiral structure in their walls (Penrose, 1994b).

'consciousness depends on biologically "orchestrated" coherent quantum processes in collections of microtubules within brain neurons'

(Hameroff and Penrose, 2014, p. 39)

in the microtubules). If quantum computing does occur in the brain, this is very important, but it only adds another layer of complexity to the way the brain works. If there is a hard problem, it might be rephrased here as ‘How does subjective experience arise from objective reduction in the microtubules?’ The strange effects entailed in quantum processes do not, of themselves, have anything to say about the experience of light or space or pain or the colour of the traffic light. They do not explain why there is experience rather than no experience. American philosopher Jesse Prinz refers to ‘those pesky quantum phenomena that are a refuge for non-mysterians with an appetite for the mysterious’ (2003, p. 116). Koch and Hepp (2006/2007; here in preprint [2007], p. 1) call the connection between quantum mechanics and consciousness an ‘entertaining topic at parties’ but argue that evidence from biophysics and computational neuroscience makes any connection unlikely. Samanta Pino and Ernesto Di Mauro (peer commentary on Hameroff and Penrose, 2014, here p. 92) wonder whether the lack of alternative explanations may be making quantum physics seem more promising than it should for tackling the ‘intricately unapproachable’ unknown that is consciousness. And Pat Churchland concludes that ‘Pixie dust in the synapses is about as explanatorily powerful as quantum coherence in the microtubules’ (1998, p. 121).

‘Pixie dust in the synapses is about as explanatorily powerful as quantum coherence in the microtubules’

(Churchland, 1998, p. 121)

But the theory may at least have the benefit of pushing us closer towards a need for ‘falsifiable verification’, Pino and Di Mauro concede (p. 92); and Hameroff and Penrose respond confidently that the proposed form of objective reduction ‘may be fairly close to either experimental confirmation or refutation’ (p. 99). So, quantum processing may yet turn out to be the solution to the hard problem, but so far it seems mainly to have moved the problem into a microtubule-shaped theatre.

MULTIPLE DRAFTS

The most concerted attempt to ditch the theatre is probably Dennett’s multiple drafts theory, which he proposed as an alternative to Cartesian materialism. ‘When you discard Cartesian dualism’, Dennett says,

you really must discard the show that would have gone on in the Cartesian Theatre, and the audience as well, for neither the show nor the audience is to be found in the brain, and the brain is the only real place there is to look for them.

(1991, p. 134)

‘Daniel Dennett is the Devil’

(Voorhees, 2000, pp. 55–56)

This wholesale rejection of the idea of ‘me’ can be deeply unsettling. So much so that Dennett has even been called ‘the Devil’, his ideas needing ‘an exorcism, aimed at eliminating the spectre of materialist reductionism from the science of consciousness’ (Voorhees, 2000, p. 55).

In multiple drafts theory, perceptions, emotions, thoughts, and all kinds of cognitive activity are accomplished in the brain by multitrack parallel processes that interpret and elaborate sensory inputs, and all are under continuous revision. Like the many drafts of a book or article, perceptions and thoughts are constantly revised and altered, and at any point in time there are multiple drafts of narrative fragments at various stages of editing in various places in the brain.

You may then want to ask, ‘but which ones are conscious?’ If you do so, you are imagining a Cartesian theatre in which only some of these drafts are re-presented

for the audience to see. If you do so, you are falling for what Dennett calls the 'myth of double transduction' (Dennett, 1998a) – the showing-again for the benefit of consciousness. This is precisely where Dennett's model differs from Cartesian materialism, for on the multiple drafts theory discriminations only have to be made once. There is no master discriminator, or self, who *has* some of the experiences. There is no 'central meander' who understands them. There are only multiple drafts all being edited at once. The sense that there is a narrative stream or sequence comes about when the parallel stream is probed in some way, such as by asking a question or requiring some kind of response. For example, some of the drafts are used in controlling actions or producing speech, and some are encoded as memory traces, while most just fade away.

A segment-like piece has been cut out of the back of his head. The whole world looks in with the sun. It makes him nervous, it distracts him from his work, and he is also irritated that he of all people should be shut out of the performance.

(Franz Kafka, diaries, 10 January 1920, our translation; for discussion see also Troscianko, 2014, pp. 106–107)

Let's suppose that you just saw a bird fly past the window. Your judgement that you consciously saw the bird is a result of probing the stream of multiple drafts at one of many possible points. There is a judgement all right, and something about the event may become accessible for future memory retrievals, but there is not also the *actual experience* of seeing the bird fly past. According to Dennett, contents arise, get revised, affect behaviour, and leave traces in memory, which then get overlaid by other traces, and so on. All this produces various narratives which are single versions of a portion of the stream of consciousness, but 'we must not make the mistake of supposing that there are facts – unrecoverable but actual facts – about just which contents were conscious and which were not at the time' (1991, p. 407). In other words, if you ask, 'what was I actually experiencing at the time the bird flew past?', there is no right answer because there is no show and no theatre in which the 'actual' experiences happen. What we come to think of as 'the fact of the matter' is the result of posthoc confabulation (see Chapter 6).

What, then, of the audience? Dennett argues that when a portion of the world (most obviously, a person, but also perhaps a computer or robot) comes to compose a skein of narratives, that portion of the world is the observer. The observer is a 'Center of Narrative Gravity' (1991, p. 410; and see Chapter 16). As contents are fixed by probing the stream at various points, as we make judgements, and as we speak about what we are doing or what we have experienced, so the benign illusion is created of there being an author. In this sense, the observer in the Cartesian theatre, real and powerful as it feels, is an illusion, even if a very special kind of illusion.

How does this theory deal with our red traffic light? If you are a Cartesian materialist, you will insist that there is some fact of the matter about whether you were or were not conscious of the light at the time. But according to the multiple drafts model, 'there are no fixed facts about the stream of consciousness independent of particular probes' (1991, p. 138), so it all depends on the way the parallel stream was probed. Had you been asked during the drive what was happening, you would

probably have noticed and then remembered the light changing to red, and been convinced that you were conscious of it. This is also why, when we ask, ‘am I conscious now?’, the answer is always yes; and why, when we ask, ‘what was I conscious of a moment ago?’, there seems to be an easy answer. But since there was no probe during your drive that led to speech or memory encoding (only those leading to changing gears and pressing pedals), you conclude – if you should happen to reflect on the journey once you get to work – that you were unconscious of the red light at the time. The only difference is because of the probes that were applied, or not. If the probing is done using language, as it often is, the role of the ‘second person’ – of social action and interaction, especially through discourse – becomes important in ways we will explore more thoroughly in Chapter 17.

‘there are no fixed facts about the stream of consciousness independent of particular probes’

(Dennett, 1991, p. 138)

We can compare this theory with the others in terms of our three crucial questions: does it help with the problem of why we have subjective experience at all, does it try to explain the difference between conscious and unconscious events, and does it posit some kind of theatre? The answer to all three is an obvious, and defiant, ‘no’. There is no theatre, no difference between conscious and unconscious events or processes, and as for subjectivity, multiple drafts theory throws out most of the assumptions that we usually make about it.



ACTIVITY 5.1

Cartesian materialism

Almost no one admits to being a Cartesian materialist, yet the literature about consciousness is full of theatrical metaphors and phrases implying that things are ‘in’ or ‘out’ of consciousness. It is worth trying to sort out what these mean before making up your own mind about the theatre of consciousness. If you are doing this as a class exercise, ask each person to find examples in advance and bring them for discussion.

Theories. For any theory of consciousness ask, ‘Does this theory use theatre imagery or metaphors? If so, is it a *Cartesian* theatre? Is this a form of Cartesian materialism?’

Tell-tale phrases. Look out for theatre imagery, or phrases that imply CM, in any area of psychology, philosophy, or neuroscience. Here are a few examples. In each case ask whether this imagery is helpful, or a sign of problems with the theory.

‘There seems to be a presence-chamber in my mind where full consciousness holds court’ (Galton, 1883, p. 203)

‘ideas [. . .] pass in rapid succession through the mind’ (James, 1890, pp. 25–26)

If we think that at any time there is a truth about what ‘I’ am subjectively experiencing now, then we are wrong. This is why Dennett is able to say, ‘But what about the *actual* phenomenology? There is no such thing’ (1991, p. 365). And this in turn may be why critics complain that Dennett has not *explained* consciousness but *explained it away*. Yet, he claims that his theory does deal with subjectivity. He describes a rich experience of sitting in his rocking chair, watching the sunlight on the trees and listening to music. This description, he says, is just one of the many possible ways the parallel stream could have been probed. If we ask, ‘but what was he actually experiencing at the time?’, there is no right answer. If we sit now and ask, ‘what am I conscious of now?’, the answer will also depend on how the stream is probed. As inner speech is produced so content becomes fixed, and we conclude that ‘I’ was watching the white fluffy clouds go by. This is how the experience and the experiencer come to be created. This is what brains do. This, according to Dennett, is how experience can be electrochemical happenings in a brain.

If you find multiple drafts difficult and worrying, then you are probably beginning to understand it. It is difficult to understand because doing so means throwing out many of our usual habits of thought concerning our own consciousness. If you want to give this theory a fair hearing before deciding on its merits, then you really need to try to understand it with an open mind, setting aside your natural assumptions. This is not easy to do, but does get easier with practice. The process may feel like identifying the routes your thoughts naturally travel down, and then gently, at a critical point, opening up a new path. And remember, if you decide, having given it a real try, that the theory doesn’t work, you can always go back to the old path.

As we have gone along we have suggested various exercises that may help to loosen up your thinking about your own consciousness – primarily asking yourself ‘Am I conscious now?’ as often as possible. Doing this will help you to assess whether Dennett’s theory really does deal with subjectivity as he claims. What is it like being you now? If Dennett is right, this question itself acts as one of many possible probes and fixes the content. Does this seem to fit with your experience?

Of course, multiple drafts theory does not solve all the problems of consciousness, or create no new problems in the attempt. If it did, this would be a book about multiple drafts, not the mystery of consciousness. Clear-sighted as Dennett can be in his criticisms of other theories, his own theory arguably runs into some of the same problems that he pinpoints in others. These arise partly from the emphasis on language, which leads to some questionable claims about what consciousness is like in the absence of ordinary human language: he dismisses the minds of deaf–mutes as ‘terribly stunted’, for instance (1991, p. 448).

Another issue is the role of the brain in multiple drafts theory. Dennett rejects the idea that a disembodied brain, or a ‘brain-in-a-vat’, could have meaningful experiences, not least because consciousness (or the illusion of consciousness) is the result of probes from outside the brain. Yet Dennett has not extended this line of thinking to take the embodied or extended mind into account, and he tends to treat the brain as the entity that does the thinking, the perceiving, the deciding: ‘the brain’, for example, ‘doesn’t always avail itself of this option’ (1991, p. 16).

The brain-centric view extends to the issue of representation. The revisions of the multiple drafts occur at the level of neural representation – by which we may infer Dennett means patterns of neural activity, or weightings of synapses, assumed to correlate with particular informational outputs. But he also talks about the phenomenology in representational terms: ‘Our visual phenomenology, the contents of visual experience, are in a format unlike any other mode of representation’ (p. 54). This can mean that the ‘content’ he talks about slips from being the content of physical representations to being the content of experiences themselves. Although he says there is no such thing as ‘the actual phenomenology’ (p. 365), he nevertheless equates ‘[o]ur visual phenomenology’ with ‘the contents of visual experience’ (p. 54), and assures us that he doesn’t mean ‘you have *no* privileged access to the nature or content of your conscious experience’ (p. 69).

A possible reply from Dennett might be to say that ‘content’ is just a metaphor, and a highly conventionalised one at that, but this is precisely the kind of metaphor which leads him to reject alternative theories. Even though sensory discriminations

‘this may help to pin down the location of awareness in the brain’ (Crick, 1994, p. 174)

‘The range and variety of conscious phenomenology [. . .] is everyman’s private theatre’ (Edelman and Tononi, 2000a, p. 20)

‘visual information that is processed in the dorsal stream does not reach conscious awareness’ (Milner, 2008, p. 195)

‘Once information is conscious, it can enter into a long series of arbitrary operations’ (Dehaene, 2014, p. 14)

Looking critically at the language people use, do you find that the concepts of the Cartesian theatre and Cartesian materialism help you distinguish between workable and unworkable theories? Or do these examples make you question the concepts? Maybe the idea of a finish line or boundary where unconscious processes turn into consciousness means that if you try to draw any distinction between conscious and unconscious you fall into the Cartesian trap, because there is no fact of the matter about whether you were conscious until you ask. On the other hand, maybe there is a way of preserving this intuitive distinction, in thinking about both the brain and your own experience, without becoming a Cartesian materialist. Or maybe you eventually conclude that it is all right to be one.

and the like do not have to be re-presented for the audience in the Cartesian theatre (p. 113), ‘collated, revised, enhanced representations’ (p. 112) sometimes seem to be what experiences either contain or are.

Some have suggested that Dennett too reimagines the theatrical show, but distributed right across the brain: ‘All the work that was dimly imagined to be done in the Cartesian Theater has to be done somewhere, and no doubt it is distributed around in the brain’ (Dennett and Kinsbourne, 1992, p. 234). As Dan Lloyd argues, ‘a specialized, localized portion of the brain’ is still required to make ‘observations’ (2000, p. 175; see Dennett, 1991, p. 113); it’s just that ‘the judgmental tasks are fragmented into many distributed moments’. Lloyd notes that early commentators ‘worried that these distributed, yet discrete, microtakings had the effect of replacing the Cartesian Theater with the Cartesian cineplex’ (Ross et al., 2000, p. 176). Can we ever really step outside the theatre, or does it wear too many disguises?

READING



Baars, B. J. (1997a). In the theatre of consciousness: Global workspace theory, a rigorous scientific theory of consciousness. *Journal of Consciousness Studies*, 4, 292–364.

A detailed debate about Baars’s theory.

Blackmore, S. (2005). *Conversations on consciousness*. New York: Oxford University Press.

Read conversations with any of the researchers discussed so far – Baars, Block, Chalmers, the Churchlands, Crick, Dennett, Hameroff, Koch, O’Regan, Penrose – and others we will meet in later chapters. This is a chance to see what people say when talking rather than writing about consciousness.

Dennett, D. C. (1991). [excerpts from several chapters]. In D. C. Dennett, *Consciousness explained* (pp. 101–115, 309–314, and 344–356). Boston, MA: Little, Brown.

Dennett’s original explanation of the Cartesian Theatre and alternatives to it.


Hameroff, S., and Penrose, R. (2014). Consciousness in the universe: A review of the ‘Orch OR’ theory. *Physics of Life Review*, 11, 39–112.

Students need not read all of the target paper, but the authors’ replies to the peer commentaries (*PLR*, 11,

94–100 and 104–112) provide helpful summaries of and responses to the criticisms (which are in *PLR*, 11, 79–93). The original article also includes sections on how the theory fits with classic ways of thinking about consciousness (p. 40) and on testable predictions from the theory (pp. 68–70).

Thomas, Nigel J.T. (2016). Mental imagery. In E. N. Zalta (Ed.), *The Stanford encyclopedia of philosophy* (Summer 2016 Edition). <http://plato.stanford.edu/archives/sum2016/entries/mental-imagery/>

Outlines the history of and current debate on the mechanisms and experiences of mental imagery, emphasising questions about representation.



The unity of consciousness

CHAPTER

SIX

'consciousness, which is itself an integral thing not made of parts, "corresponds" to the entire activity of the brain'

(James, 1890, i, p. 177)

Why do we seem to have only *one* consciousness? Why is consciousness, as James puts it, 'an integral thing not made of parts' (1890, i, p. 177), when the brain is such a massively parallel, and complicated, multi-tasking organ? Why do we feel as though there is just one conscious mind in here that is experiencing a unified world? This classic problem from philosophy takes on new significance in the context of modern neuroscience (Cleeremans, 2003).

The problem is simple to state. When we look at the brain side of the great divide, we see nothing but complexity and diversity. At any given time, countless different processes are all going on at once, and in different areas. For example, right now your visual system is processing multiple inputs and dealing with colour, motion, and other features in different areas of the eyes and brain. At the same time, processing is going on in other sensory areas, in memory systems, and in emotional systems. Thoughts are bubbling along, movements are being planned and coordinated, and sentences are being constructed. All these diverse processes are linked through multiple routes and connections, but, as we have already seen (Chapter 5), there is no single place in the brain where everything is brought together for someone to watch. All the parts just keep doing their things all at once.

The parts also take time to do their things. Stuff happens quickly in the brain, but not instantaneously, and not all at the same speed. Signals travel along neurons at about 100 m per second, but this varies depending on the width and myelination of a given neuron, ranging between about 0.5 m per second along the small-diameter, unmyelinated pain receptors, and up to 120 m per second in the large myelinated neurons linking the spinal cord to the muscles. Signals also take time

to jump between neurons at the small gaps between them, the synapses. Crossing a synapse takes at least half a millisecond, so the more neurons are involved in a given process, the longer that process takes (Welsh, 2015).

But when we look at the mind side of the divide, everything seems to be unified. It seems as though, right now, there is one 'me' and one more or less continuous stream of experiences happening to me now. The German philosopher Thomas Metzinger claims that experience requires there to be unity in what is experienced: 'In order for a world to appear to us, it has to be *one* world first' (2009, p. 27). He also says that unity can only ever be experienced as temporally *present*: 'One thing cannot be doubted from the first-person perspective: I always experience the wholeness of reality *now*' (1995a, p. 429). The question is, how can such unity, experienced right now, arise from such diversity of non-instantaneous processing?

We will tackle this question from a variety of perspectives, leaving for later the important questions about the unity of self (Chapter 16). We will explore how the different features of objects are brought together to make a single object (the binding problem), and how the different senses are brought together to make a unified experienced world (multisensory integration). We will investigate how subjective and clock time are integrated with each other. Finally, we will consider what happens when consciousness is more or less unified than normal, using examples from synaesthesia, split brains, amnesia, and neglect. These atypical cases may make us question what we assume about normal experience.

Among the many theories about unity, a tempting but probably unworkable option is dualism. Substance dualists mostly believe that consciousness is intrinsically unitary, each person having their own single consciousness distinct from their physical brain. Indeed, it was partly the argument from unity that led Descartes to his dualism. It was also the argument from unity that led to Popper and Eccles to their 'dualist interactionism'. Their preferred solution was that '*the unity of conscious experience is provided by the self-conscious mind and not by the neural machinery of the liaison areas of the cerebral hemisphere*' (1977, p. 362). They argued that the mind plays an active role in selecting, reading out, and integrating neural activity, moulding it into a unified whole according to its desire or interest. The problem for Popper and Eccles, as for all dualists, is how this mind–brain interaction takes place. The theory provides no explanation of *how* the separate mind carries out its selecting and unifying tasks, and for this reason very few people accept it.

Another dualist, though not a substance dualist, was Benjamin Libet, who believed that conscious unity was achieved through the effects of a 'conscious mental field' (CMF). Libet was a scientist unafraid of putting his ideas to the empirical test, and proposed the following experiment: take an isolated piece of cortical tissue that is completely cut off from the rest of the brain but kept fully functioning and alive, then activate it electrically or chemically. If there is a CMF, this stimulation should produce a conscious experience in the person who has the rest of the brain. 'Communication would then have to take place in the form of some field that does not depend on nerve pathways' (Libet, 2004, p. 172). This sounds like a form of telepathy within one brain, and presumably most scientists would expect the experiment to fail. Nevertheless, it could be tested.

'The unity of consciousness is illusory'

(Hilgard, 1977, p. 1)

'I cannot distinguish in myself any parts, but apprehend myself to be clearly one and entire'

(Descartes, 1641/1970, p. 196)



PRACTICE 6.1

IS THIS EXPERIENCE UNIFIED?

As many times as you can every day, ask yourself, '***Is this experience unified?***'

You might like to begin, as usual, by asking, 'Am I conscious now?' and then explore what you are conscious of, all the time attending to whether the experience is unified. You could try this: pay attention to your visual experience for a few seconds. Now switch to sounds. You will probably be aware of sounds that have been going on for some time. Has the sight just become unified with the sound? What was going on before? What role does attention play in this? You can do the same with verbal thoughts and bodily sensations. Is your consciousness always unified? Is it now?

Since materialism encounters obvious problems too, there are still researchers who defend versions of mind–body dualism and, despite its 'bad reputation', defend it as part of a 'progressive research programme' for exploring all kinds of questions about consciousness, including its apparent unity (Lavazza and Robinson, 2014, pp. 7, 5). But a more common approach is to try to find out how the brain and the rest of the body manage to integrate and unify their functions, and the majority of the examples considered here attempt this. A third and final approach is to reject the idea that consciousness really is unified at all. Perhaps, on closer inspection, we might find that the apparent unity is illusory. In this case, the task is to explain how we can be so deluded.

*'I am one person living
in one world'*

(Metzinger, 1995b, p. 427)

THE BINDING PROBLEM

Take a coin, toss it, and catch it again in your hand. How does this object, the coin, appear to you as it flies? You might like to toss it a few times and watch carefully. What do you see?

You will probably see a single object fly up in the air, twist over, and land in one piece on your hand. Bits don't fly off. The silver colour doesn't depart from the shape, and the shape doesn't lag behind the motion. But why not?

Think now of what is happening in the visual system. Information extracted from a rich and rapidly changing pattern of excitation in the rods and cones of the retina takes one route through the superior colliculus to the eye-movement system and thereby controls your visual tracking of the moving coin. Other information from the same retinal patterns takes a different route through the lateral geniculate nucleus to visual cortex. In V1 there are many retinotopic maps (that is, the organisation of cells reflects the layout on the retina), and here the hierarchical processing of edges, lines, and other basic features begins. Meanwhile, other visual cortical areas are handling other features, including colour, movement,

and shape or form. In these higher areas, the original mapping is lost, and features are dealt with regardless of where on the retina they originally fell.

These different processes take different lengths of time. For example, colour is processed faster than orientation, and orientation faster than motion. Then there are the two major visual streams to consider: the dorsal stream controlling the fast action of catching the coin deftly (if you did), and the ventral stream involved in the more time-consuming process of perceiving the coin as a coin, with the two in complex dynamic interaction with each other. There is no single place and time in the brain at which everything comes together for the falling coin to be consciously perceived as one thing rather than a floating collection of attributes. How, then, do we consciously perceive a falling coin as one moving object?

The problem described here is that of visual binding, but the more general 'binding problem' applies across many different sensory modalities, and at many levels of description, from the neural level to the phenomenological (Bayne and Chalmers, 2003). Some of these levels (most obviously the neural processes) can be studied quite separately from consciousness (Revonsuo, 1999, 2009). Towards the more cognitive end, the problem is how conjunctions of features are represented, ranging from the binding of shape and colour in detecting blue triangles or red squares, to the binding of words and phrases with their roles in sentences. In particular, the problem for consciousness is how this kind of binding happens dynamically in real time. As the coin flips, what keeps the colour, form, movement, and other attributes of the coin together?

This problem is intimately bound up with both memory and attention. For example, try remembering entering your own front door. To do this successfully, various features have to be imagined at once: maybe the colour of the door, the flowers growing round it, or the pile of rubbish in the corner, probably the key on its ring and the way you have to turn it. We considered in Chapter 3 the question of just how detailed this experience really is; if vision does not actually operate by building up pictures of the world, maybe lots of this information is simply not given. For example, when I imagine a face, my imaginative experience may not specify whether or not the person is wearing glasses (Pylyshyn, 2003, p. 34). And when you imagine your door, the pile of rubbish may be neither present nor explicitly absent; your imagining may simply not specify. But the concept of binding does not commit us to any particular position as regards the amount of detail; even if all we imagine is turning a key in a nondescript door, there are still things that need binding to each other. And as a result of the binding, you experience a more or less unified memory of something you do every day. All that information is held briefly in working memory and, as we have seen, some theories of consciousness, such as global workspace theories, relate consciousness to working memory, and to attentional amplification of fronto-parietal circuits.

Some people argue that the binding problem is precisely the same problem as understanding how attention works – a topic to which we will return in the next




FIGURE 6.1 • Is there a flipping coin somewhere in the mind or brain? We know that the colour, motion, and shape of the coin are processed in different brain areas, but how are these features bound together to produce the single experience of a flipping coin? The binding problem cannot be solved by imagining that they are all brought together for display to an observer inside the brain.

chapter. On this view, as long as you pay attention to the flipping coin or the image of the door, their various attributes are bound together. When you think about something else, the diverse attributes fall apart, and the coin or door is no longer experienced as a unified whole.

There is evidence that attention is required for binding. For example, when people's attention is overloaded or diverted, the wrong features can be bound together to produce illusory conjunctions, such as when you are rushing along the street and see a black dog, only to realise that it was in fact a golden Labrador passing a black rubbish bag.

Bilateral damage to parietal cortex, which affects attention, can cause binding deficits, and in visual search tasks focused attention is necessary for finding unknown conjunctions. Anne Treisman, a British psychologist based at Princeton, interprets the relationship in terms of 'feature integration theory' (Treisman and Gelade, 1980). When we attend to objects, computationally understood 'temporary object files' bind groups of features together on the basis of their spatial locations. For Treisman (2003), binding is central to conscious experience, and conscious access in perception is always to bound objects and events, not to free-floating features of those objects or events (see also Merker, 2013).

Other factors suggest that, closely related as they are, binding and attention cannot be the same thing. Think of how you caught the coin. The fast visuomotor control system in the dorsal stream has a complex computational task to carry out in real time. It must track the current speed and trajectory of the coin and direct your hand, with the right fingers in position, to catch the coin as it drops. For this task, the form and movement of the coin must be bound together with each other and not with the movement of some other object in the vicinity. If you swat away a fly, return a fast serve, or avoid a puddle as you run down the street, the features of these objects must be well enough bound together to be treated as wholes. Yet, as we will explore further in the next two chapters, you do all these things very fast and often without paying attention. There are obviously close relationships between attention, consciousness, and binding, but just what sort of relationships is not yet clear.

BINDING BY SYNCHRONY

The best-known theory relating binding and consciousness is that proposed by Francis Crick and Christof Koch (1990). In the 1980s, studies of the cat's visual cortex had revealed oscillations in the range of 35 to 75 hertz (i.e. 35–75 cycles per second), in which large numbers of neurons all fired in synchrony. These are often referred to as 'gamma oscillations' or (rather inaccurately) as 40-hertz oscillations, and the idea was that all the neurons dealing with attributes of a single object would bind these attributes together by firing in synchrony. According to Crick and Koch, 'this synchronized firing on, or near, the beat of a gamma oscillation (in the 35- to 75-hertz range) *might be the neural correlate of visual awareness*' (Crick, 1994, p. 245).

They argued that the thalamus controls attention by selecting the features to be bound together by synchronisation of firing. The Columbian physiologist Rodolfo Llinás used a similar idea to account for temporal binding and ultimately the unity of self, arguing that 'consciousness is a product of thalamocortical activity' (2002, p. 131).

Crick concluded that ‘Consciousness [. . .] exists only if certain cortical areas have reverberatory circuits (involving cortical layers 4 and 6) that project strongly enough to produce significant reverberations’ (Crick, 1994, p. 252). He argued that it may be more efficient for the brain to have ‘one single explicit representation’ rather than sending tacit information to many different parts of the brain. In other words, he distinguished explicit and tacit (conscious and unconscious) information, and thought that the unity of consciousness is real and not illusory.

In their later work, Crick and Koch (2003) gave up the idea that 40-hertz oscillations are a sufficient condition for the NCC, arguing instead that the features of a single object or event are bound together when they form part of one temporary coalition of neurons, and that the primary role of synchrony is to help one coalition in the competition for consciousness. EEG studies have confirmed a role for synchrony in visual binding (Tallon-Baudry and Bertrand, 1999; Tallon-Baudry, 2003). Catherine Tallon-Baudry and her colleagues in Lyon showed that gamma oscillations were much stronger during feature-binding tasks, and for tasks in which participants had to hold the representation of an object in short-term memory while searching a display, than during control tasks. They also studied two patients who had electrodes implanted in extrastriate visual areas during their treatment for epilepsy. When these patients were asked to keep a visual image in mind while performing a matching task, visual areas separated by several centimetres became synchronised, with oscillations in the beta range (15–25 Hz). They took these results to confirm Hebb’s (1949) fifty-year-old suggestion that short-term memory is sustained by reverberating activity in neuronal loops.

Another experiment used a modified version of the famous image of a Dalmatian dog hidden in a black and white meaningless pattern. Seeing the dog correlated with an increased EEG response in the gamma band. From these and other studies, Tallon-Baudry concluded that any stimulus elicits locally synchronised activity in early visual areas, sufficient for coarse and unconscious identification. These local oscillations could then be more strongly synchronised between areas ‘to provide a much more detailed representation of the stimulus, along maybe with the conscious experience of it’ (2003, p. 361). As so often with neuroscientific studies, however, consciousness is brought into the discussion only as an added extra at the very end.

Binding by synchrony does not necessarily involve gamma oscillations. In their model of temporal binding, Andreas Engel, Wolf Singer, and their colleagues in Frankfurt, Germany, propose that objects are represented in the visual cortex by assemblies of synchronously firing neurons (Engel et al., 1999; Engel, 2003; Singer, 2000, 2007). For example, in Figure 6.3 the lady and her cat are each represented by one such assembly of cells. Each assembly consists of neurons which detect

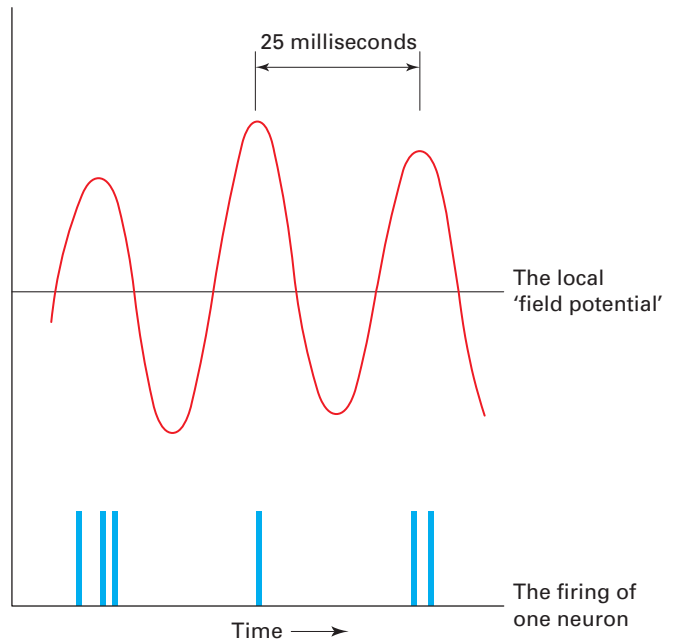


FIGURE 6.2 • A simple figure to illustrate how neurons fire in a 40-hertz rhythm. (A 40-hertz oscillation repeats every 25 msec.) The smooth curve represents the local field potential. This is a measure of the average ‘activity’ of many of the neurons in that neighbourhood. The short vertical lines show the firing of just one neuron. Notice how, when this neuron fires, it fires ‘on the beat’ of some of its neighbours, represented by the local field potential. The usual sign convention for plotting the field potentials has been reversed (Crick, 1994, p. 245).




FIGURE 6.3 • Engel and colleagues' model of temporal binding. The model assumes establishment of coherent representational states by temporal binding. The model assumes that objects are represented in the visual cortex by assemblies of synchronously firing neurons. In this example, the lady and her cat would each be represented by one such assembly (indicated by open and filled symbols, respectively). These assemblies comprise neurons which detect specific features of visual objects (such as, for instance, the orientation of contour segments) within their receptive fields (lower left). The relationship between the features can then be encoded by the temporal correlation among these neurons (lower right). The model assumes that neurons which are part of the same assembly fire in synchrony, whereas no consistent temporal relation is found between cells belonging to different object representations (from Engel et al., 1999, p. 131).

specific features of the objects, such as colour, movement, or orientation. The relationship between these features is then encoded by temporal correlation among the neurons. This means that neurons representing one object fire in synchrony with each other, but out of synchrony with neurons representing other objects at the same time. This means figure and ground can be segregated and individual objects represented without confusing them. The model allows for many kinds of synchronised neural activity apart from that based on oscillations.

Engel and his colleagues reviewed many studies showing that synchronisation between cells occurs widely in both perceptual and motor systems. They conclude that arousal and selective attention involve enhanced synchrony in the relevant populations of neurons, and that 'temporal

binding may indeed be a prerequisite for the access of information to phenomenal consciousness' (Engel et al., 1999, p. 133). In their view, synchronisation is necessary for consciousness but not sufficient because information must also enter short-term memory. This brings the theory closer to GWT.

The theoretical and empirical basis of binding by synchrony remains controversial. Some experiments have failed to find any relationship between synchrony and binding of moving patterns or static visual objects (e.g. Thiele and Stoner, 2003; Dong et al., 2008). Some researchers have also argued that the theory is incoherent in being focused on early stages of cortical processing, even though both neurological evidence and the perceptual facts of binding suggest that it must be a high-level computation; and that the architecture of the cerebral cortex lacks the mechanisms needed to decode synchronous spikes and to treat them as a special code (Shadlen and Movshon, 1999). Others have even suggested that none of this is needed because the 'binding problem' does not exist (Merker, 2013).

These attempts to solve the binding problem seem related to the unity of consciousness, but although they explain how percepts are unified, we may still wonder how this relates to subjectivity. Some conclusions, like that of Engel and colleagues above, are ambiguous about correlation and causation, for example when they conclude that 'at least at early stages of sensory processing, the degree of synchronicity *predicts* reliably whether neural activity will contribute to conscious experience or not' (Engel et al., 1999, p. 146, our italics). Some of their phrases also imply Cartesian materialism – including talk of information having access to phenomenal consciousness (e.g. 1999, pp. 133, 141, 144). Although they are trying to explain the unity of phenomenal awareness (subjectivity, or 'what it's like') in neural terms and without magic, they seem to imply that information is

unified first and then somehow ‘enters consciousness.’ What this means remains unexplained.

Proceeding from the ‘close relation between consciousness and binding,’ Singer suggests that ‘only those results of the numerous computational processes that have been bound successfully will enter consciousness simultaneously’ (in Metzinger, 2009, p. 67). Although ‘entering consciousness’ sounds like magic, he explains that consciousness does not depend on any particular group of neurons but is ‘an emergent property of a particular dynamical state of the distributed cortical network.’ Yet he does not explain how this emergence works, nor what it means for subjectivity to emerge from the temporal coherence of a large population of distributed neurons.

*‘only those results
[...] that have been
bound successfully will
enter consciousness
simultaneously’*

(Singer, in Metzinger, 2009, p. 67)

MICRO-CONSCIOUSNESSES

The very concept of the unity of consciousness is questioned by British neurobiologist Semir Zeki, who proposes instead that there are as many micro-consciousnesses as processing nodes in a system. He is not referring to cases of multiple personality or split brains (considered later in this chapter), but claiming that a multiplicity of consciousnesses is the norm: the unification that comes with self-consciousness is an exception made possible only by language.

Zeki (2001, 2007) notes that the many parallel systems of the primate visual system reach their perceptual endpoints at different times. For example, the attributes of an object, such as orientation, depth, facial expression, or shape, are processed separately, and colour (processed in V4 and V4a) may be as much as 80 ms ahead of motion (processed in V5) (Zeki and Bartels, 1999). But we are not aware of this perceptual asynchrony. In Zeki’s view, each of these separate cortical systems has its own conscious correlate.

*‘binding is a post-
conscious phenomenon’*

(Zeki, 2007, p. 584)

Cortical and subcortical systems both consist of many hierarchically organised nodes with multiple inputs, outputs, and connections, and since no node is a recipient only, ‘there is no terminal station in the cortex’ (Zeki, 2001, pp. 60–61), ‘no final integrator station in the brain’ (Zeki and Bartels, 1999, p. 225), no ‘pontifical neuron’ (James, 1890). There is no need for micro-consciousnesses to ‘be reported to a “center” for consciousness, or a “Cartesian Theater”’ (Zeki, 2001, p. 69). ‘Visual consciousness is therefore distributed in both space and time’ (p. 57).

Experiments inspired by Mondrian’s abstract artworks show that basic colour perception can occur without the activation of the frontal lobes (Zeki and Marini, 1998), which implies phenomenal consciousness without higher processing. Conversely, in patients with lesions that cut off input from primary visual cortex, activity in certain cells of V5 (which plays a major role in motion perception) can correlate with ‘a crude but conscious experience’ (Zeki, 2015, p. 12) of fast high-contrast movement. This indicates that active participation by V1 is not required for visual consciousness, as had long been believed.

Because micro-consciousnesses are distributed in both space and time, and because binding of different attributes takes different lengths of time, Zeki argues that binding is a post-conscious phenomenon (Zeki and Bartels 1998, 1999; Zeki, 2007). In other words, there is phenomenal consciousness of visual attributes before those attributes are bound. This sets Zeki apart from Crick and Koch (1990;

IS THIS EXPERIENCE UNIFIED?

Koch, 2004), who argue that consciousness is built up only when stable coalitions form; from Engel and Singer (above); and from Metzinger, who argues that ‘Consciousness is what binds things together into a comprehensive, simultaneous whole’, and that only if the senses are unified can you experience a world (2009, p. 26).

Zeki’s suggestion is not a form of panpsychism (the view that everything is conscious), because he does not claim that *all* neural processing is conscious. This means that, unlike Dennett, he must still distinguish between conscious and unconscious processes. He speculates that neural activity remains implicit as long as it requires further processing; when processing is complete, it becomes explicit or conscious. Discussing blindsight (Chapter 8) and other kinds of brain damage, he describes people who report being conscious of aspects of early motion processing that would be implicit in other people, leading him to propose that ‘cells whose activity is only implicit can, in the right circumstances, become explicit. Put more boldly, cells can have double duties, rendering the incoming signals explicit or not, depending on the activity at the next node’ (2001, p. 66).

In Zeki’s (2007) scheme, micro-consciousnesses are bound into macro-consciousnesses which correspond to Block’s phenomenal consciousness. Unified consciousness corresponds to Block’s access consciousness and comes about only through language and the awareness of a perceiving self. This is very much a ‘consciousness at the top’ model as discussed in Chapter 5, but one with the idiosyncrasy of including other consciousnesses ‘lower down’. These multiple consciousnesses constitute a hierarchy ‘with what Kant called the “synthetic, transcendental” unified consciousness (that of myself as the perceiving person) sitting at the apex’ (Zeki, 2003, p. 214). Elsewhere, however, Zeki remarks that ‘the term microconsciousness may itself be a misnomer, because it implicitly supposes that there is a higher unified and singular conscious entity, beyond all the microconsciousnesses’ (Zeki and Bartels, 1999, p. 238).

One way of interpreting Zeki’s micro-consciousnesses is that whether bound or not, they are all conscious, implying a head full of disconnected phenomenal experiences along with a unified world perceived by a constructed self. This would avoid the problem of explaining the difference between conscious and unconscious processes. But this cannot be what he means, for he says that ‘once a macro-consciousness is formed from two or more micro-consciousnesses, the constituent micro-consciousnesses cease to exist’ (2007, p. 584), and we become aware of the composite instead – though by paying attention to the constituents rather than the composite we can reverse this effect. The tricky point for his theory is the point in space or time at which information that was previously ‘implicit’ is ‘rendered explicit’ or ‘acquires a conscious correlate’. This transition remains, as in every other theory that comes up against it, unexplained and essentially mysterious.

MULTISENSORY INTEGRATION

The theories considered so far have concentrated on vision, perhaps because vision is the predominant sense in humans, and also because the visual system is much better understood than other parts of the brain. But binding has to occur between the senses as well as within them. This is no trivial matter, not least

because the senses are so different. For example, while vision depends largely on spatial analysis, hearing uses temporal analysis. How are these two very different processes integrated?

Think of the way you turn your head and eyes to look straight at someone who calls your name, or the way that the smell and touch and sight of the sandwich in your hand all seem to belong to the same object. Or think of a cat out hunting. It listens to the rustling in the undergrowth, creeps carefully between the leaves, feeling its way with its whiskers, then spies the poor vole, and pounces. Cats and rats construct a spatial map around themselves in which information from their eyes, ears, whiskers, and paws is all integrated. When information from both sight and sound comes from the same position, then responses are enhanced, and information from one sense can affect responses to another sense in many ways (Stein, Wallace, and Stanford, 2001).

Somehow decisions are made about what to integrate and when. Principles of time, space, and inverse effectiveness have been proposed as guiding the integration process: if cross-modal stimuli arise at the same time and place, and if in isolation these stimuli evoke relatively weak responses, multisensory integration is more likely, and likely to be stronger. A more general framework for understanding the likelihood of different levels of segregation and integration in multisensory perception is the Bayesian model, in which new evidence is combined with prior belief to assess probability (Beierholm, Quartz, and Shams, 2009).

Within the brain, integration depends on multisensory neurons that respond to input from more than one sense. In the superior colliculus in the midbrain, cells may respond to more than one sense even at birth, but their multisensory capacity increases with experience and with increasing connections into the cortex.

Integration can give rise to illusions. Ventriloquism (Recanzone, 2009) works because when you hear a voice speaking words that coincide with movements of a toy's mouth, you hear the sounds as though they come from the toy. Experiments with the McGurk effect entail watching a person speaking one phoneme while listening to a different one, and the result can be a different phoneme altogether. For example, if you are shown someone saying 'ga' and played the sound 'ba', you will hear 'da'.

Nonetheless, for most of us, most of the time, the senses remain easily distinguishable. That is, we are not confused as to whether we heard, saw, or touched something. This ability is not as obvious as it may seem, because all the senses work by using the same kinds of neural impulses. So it seems that some explanation is needed for why they are experienced as so distinct (O'Regan and Noë, 2001). Perhaps we can learn something from the phenomenon of synaesthesia, in which the senses are not so distinct (see Concept 6.1). Many people can remember that as children they sometimes heard smells or tasted sounds, and in some people this mixing of the senses remains part of their lifelong experience.

Some argue that the more we learn about interaction between senses, the harder it is to define what a 'sensory modality' means: do pain, vestibular awareness, or thermal perception count as separate modalities, and if not, why not? And what about awareness of speech or music compared to general sound perception, or the interactions between apparently distinct



PROFILE 6.1

Giulio Tononi



Giulio Tononi is a neuroscientist and psychiatrist based at the University of Wisconsin, where he holds chairs in sleep medicine as well as in consciousness science. After studying medicine at the University of Pisa in Italy, he specialised in psychiatry and served as a medical officer in

the army before doing a doctorate in neuroscience. Long fascinated by sleep and the reasons we need so much of it, he has worked on human, mouse, and fruit-fly models; explored genetics, proteins, and computer analysis; and, with Chiara Cirelli, developed the 'synaptic homeostasis' hypothesis that sleep serves to regulate the excessive synaptic activation of wakefulness. Together with Gerald Edelman, he developed the dynamic core hypothesis, a model of consciousness which he expanded into integrated information theory and has continued to update. In IIT, a system's consciousness is determined by its causal properties and corresponds to a system's capacity to integrate information, an idea which is supported by the breakdown of information integration in slow-wave sleep, general anaesthesia, and vegetative states.

senses in sensory substitution (Chapter 8)? On the 'moderate sensory pluralism' account (Fulkerson, 2014), individual senses and the multimodal interactions between them have to be differently categorised depending on the context, so puzzling over the integration of discrete modalities may be creating a problem where there isn't one.

Some combination of sensory distinctness and multisensory integration makes possible a world in which objects can be recognised as whole, and as being touched, seen, tasted, smelled, or heard, and as being the same thing however we perceive them. Many brain areas are known to be involved, including primary sensory cortices, frontal lobe, the superior colliculus in the midbrain, and many subcortical areas, but just how this kind of integration gives rise to the subjective sense of being one self in a unified world remains to be seen.

INTEGRATED INFORMATION THEORY

The unity of consciousness is the starting point for the integrated information theory (IIT) of consciousness (Tononi, 2004, 2007, 2008, 2015), discussed briefly in Chapter 5. This developed out of Edelman and Tononi's (2000a, 2000b) dynamic core hypothesis, in which re-entrant thalamocortical loops produce high levels of dynamic complexity that create consciousness. According to IIT, consciousness corresponds to the capacity of a system to integrate information. Information is *integrated* if it cannot be localised in any individual part of the system, or is 'generated by causal interactions in the whole, over and above the information generated by the parts' (Tononi, 2008, p. 221).

IIT sets out to explain five key features of consciousness.

- 1 **Intrinsic existence.** My conscious experience exists here and now, from my intrinsic perspective.
- 2 **Composition.** Consciousness is structured: it is composed of multiple phenomenal distinctions at different levels of generality, like a red colour, a traffic light, the left side, a red traffic light on the left, etc.
- 3 **Information.** Every conscious experience is specific and differentiated from other possible experiences; its distinctions specify spatial locations, as well as positive concepts – like road as opposed to no road, red as opposed to amber or green, etc. – and negative concepts – no red, no country lane, no beach, etc.
- 4 **Exclusion.** Each experience is definite in both content and spatio-temporal 'grain', making some phenomenal distinctions – driving along a road and coming to a red light – and not others, and flowing at one speed not another.
- 5 **Integration.** Seeing the red traffic light cannot be reduced to seeing the colour red plus a traffic light.

Unlike illusionist accounts of consciousness, IIT is based on the idea that we can have complete confidence about our own conscious experiences: ‘the existence of one’s consciousness and its other essential properties is certain’, Tononi (2015) claims, and from there the theory proceeds to make inferences about the causes of consciousness from that supposedly secure basis.

‘Consciousness is unified: it is irreducible to non-interdependent, disjoint subsets of phenomenal distinctions’, says Tononi (2015). At the same time, it is extraordinarily informative, appearing to contain countless infinitesimally small chunks of information. Yet what makes a conscious state informative is not how much information it may or may not contain, but the fact that it is just one of potentially billions of other possible states. Consider a human and a photodiode facing a blank screen that is alternatively on or off. The photodiode can make only two distinctions: ‘light’ or ‘dark’. The human can distinguish the light screen not just from the dark screen, but also from a red and a green screen, and from other screens, showing any number of films, and from sounds, thoughts, and so on. This is a vast amount of information, for you can discriminate between all these states and each state has different behavioural consequences. So, you are vastly more conscious than the photodiode.

The theory proposes that each of the five essential properties of experience must be accounted for by a corresponding causal property of the physical system – the brain. The axiom of integration, or unity, means that every part of the system must be able to both affect and be affected by every other part, because otherwise the integration would be reducible to a subsection of the system. According to IIT, ‘experience is a maximum of intrinsically irreducible cause–effect power’ (Tononi, 2015). This irreducibility is measured as integrated information, referred to in the theory as ϕ , Φ , and this is calculated according to a series of mathematical formulae which you can find in Tononi’s 2015 description. This means that consciousness can be graded rather than being all-or-none. The specific anatomical location of the neural substrate of consciousness is not yet specified by the theory, but Tononi states that whether it turns out to be distributed among most cortical areas or only a subset of them, and whether it includes all cortical layers or only particular cell types, ‘IIT predicts that in each case the neural substrate of consciousness should be a local maximum of information integration’ (Tononi, 2015).

Tononi suggests that IIT allows the possibility of zombies because there could be systems which look identical to humans from an external perspective, but whose physical substrate consists of lots of mini-complexes of a low maximum Φ value rather than forming a large complex of high maximum Φ . Physical transistors in a computer are not like neurons because they cannot be grouped into macro-elements with irreducible structures. ‘Hence the brain is conscious and the




FIGURE 6.4 • Axioms of Integrated Information Theory (Tononi, 2015). The illustration is a colourised version of Ernst Mach’s ‘View from the left eye’.

‘consciousness varies with integrated information’

(Tononi and Koch, 2015, p. 15)

computer is not – it would have zero Φ and be a perfect zombie’ (Tononi, 2015). Yet surely this system would not behave identically to a human, and so in that important sense would not fit the definition of a zombie – let alone be a perfect one.

Like global workspace models, IIT insists on the importance of distributed dynamic processes and treats consciousness as a continuous variable. But in other respects, the theories differ. In GWTs, the contents of the workspace are conscious because they are displayed or made available to the rest of the system. In IIT there is no equivalent of this theatre-like display or global availability, other than the distributed power to affect other parts of the brain. Experience is a fundamental quantity, like mass, charge, or energy (Tononi, 2004), and consciousness ‘increases in proportion to a system’s ability to integrate information’ (Tononi, 2007, p. 298).

A great deal rests on the central concept of integrated information, Φ . Although there are several competing definitions of the value of Φ , the most common version seems to be, essentially, that the value of Φ is obtained by dividing the system into parts A and B and minimising a measure of the shared information between A’s outputs and B’s inputs and vice versa. According to the theoretical computer scientist Scott Aaronson (2014), however, having a large Φ value cannot be a sufficient condition for consciousness. Like Dennett in several contexts we have already come across, Aaronson (2014) warns us that our intuition about the magical status of integrated complexity may be leading us astray:

As humans, we seem to have the intuition that global integration of information is such a powerful property that no ‘simple’ or ‘mundane’ computational process could possibly achieve it. But our intuition is wrong. If it were right, then we wouldn’t have linear-size superconcentrators or LDPC [low-density parity check] codes.

A superconcentrator is a type of graph that appears in the design of communication networks, and a parity check code is an error-checking code that ensures correct transmission of a digital signal. Both rely on highly complex maths and have huge expressive power, but other than their integration of lots of information, they don’t offer very persuasive reasons to consider them conscious.

Aaronson proposes a slightly easier version of the hard problem, the Pretty-Hard Problem of Consciousness: ‘which physical systems are associated with consciousness and which aren’t’. In his view, IIT cannot even solve the Pretty-Hard Problem, let alone Chalmers’ *Hard* Hard problem, ‘because it unavoidably predicts vast amounts of consciousness in physical systems that no sane person would regard as particularly “conscious” at all’. As he puts it, ‘you can have integrated information without consciousness (or even intelligence) – just like you can have computation without consciousness, and unpredictability without consciousness, and electricity without consciousness’ (2014). But of course this just pits one intuition against another. When using a theory to make predictions about which non-human systems are conscious, it is hard to know whether the theory or your intuitions should win out. If your theory gives a fridge and a paving stone high enough scores to make them conscious, should you reject the theory or accept its predictions? In Chapter 12, we will explore the problems of testing for consciousness in human-made machines.

‘IIT unavoidably predicts vast amounts of consciousness in physical systems that no sane person would regard as particularly “conscious” at all’

(Aaronson, 2014)

Commenting on Aaronson's blog post (comment #125), Chalmers splits the Pretty-Hard Problem up into four subversions, and suggests that IIT may still be a candidate partial answer to a version in which we try to match the facts rather than our intuitions about which systems are conscious. On present evidence, it doesn't seem that IIT could ever offer more than a partial answer to a partial question, but it does have the advantage of entailing specific testable hypotheses, both mathematically and empirically. Maybe this alone justifies its current popularity.

UNITY IN ACTION

Theories discussed so far either leave the hard problem untouched or involve magical transformations of neural firing into subjective experience. Escaping completely from these problems is very difficult. One way forward might be to drop the idea of unifying representations or experiences, and think instead of unity of action. British biophysicist Rodney Cotterill says:

I believe that the problem confronted during evolution of complex organisms like ourselves was not to unify conscious experience but rather to avoid destroying the unity that Nature provided. [. . .] singleness of action is a vital requirement; if motor responses were not unified, an animal could quite literally tear itself apart!

(1995, p. 301)

He concludes that consciousness arises through an interaction between brain, body, and environment.

British philosopher Susan Hurley (1954–2007) rejects the conventional idea of consciousness as a filling in the 'Classical Sandwich' between input and output or perception and action (Hurley, 1998). Instead, she stresses that perception, action, and environment are intimately intertwined. The unity of consciousness arises from a dynamic stream of low-level causal processes and multiple feedback loops linking input and output, in an organism which she describes as a loosely centred 'dynamic singularity' with no clear external boundaries.

'unity [is] more like the twisting together of the strands of a rope, where each strand displays continuity of sensory and motor aspects'

In a similar vein, Nicholas Humphrey asks what makes the parts of a person belong together – if and when they do. Although Humphrey himself may be made up of many different selves, he concludes that

(Hurley, 1998, p. 183)

these selves have come to belong together as the one Self that I am because they are engaged in one and the same enterprise: the enterprise of steering me – body and soul – through the physical and social world. [. . .] my selves have become co-conscious through collaboration.

(2002, p. 12)

These views are all versions of enactive theories of consciousness. They treat consciousness as a kind of acting or doing, rather than representing or receiving information, so being conscious means interacting with the world, or reaching out to the world. This sidesteps the question of whether consciousness is really unified or not, for it is obvious that a single organism, whether an amoeba or a woman, has to have unified action. The tricky part is to understand how acting,

even in a unified way, can *feel* like something. This is tackled most directly by sensorimotor theory (O'Regan and Noë, 2001, Chapter 3), in which the *feeling* of sensations comes about while we are acting; being conscious is actively mastering the contingencies between the external world and what we can do with it. *What it's like* is not something that has to be mysteriously generated; it is naturally constituted by the fact of being engaged in exercising a sensorimotor skill, and specific *what-it's-likes* can be characterised in terms of dimensions like richness (how much information is available), bodiliness (how bodily movements cause sensory changes), insubordinateness (how sensory input changes without the observer's voluntary control), and grabbiness (how events grab our cognitive systems). Like IIT, this model tries to capture why sensory experiences feel different from each other, but also why each one feels unified – in this case as part of embodied action (O'Regan, 2011, e.g. p. 165).

UNITY AS ILLUSION

Finally, some people reject the notion that consciousness is unified at all. James asked of consciousness, 'does it only seem continuous to itself by an illusion?' (1890, i, p. 200). We have questioned whether the stream of conscious vision could be an illusion. Could the apparent unity of consciousness be an illusion too? This question is complicated by the fact that whenever we ask ourselves 'am I conscious now?', the answer always seems to be 'yes'. We cannot catch ourselves out *not* being conscious, and when we do find ourselves being conscious there seems to be one me and one unified experience. But what is it like the rest of the time?

One possibility is that there is nothing it is like *for me* most of the time (Blackmore, 2002, 2011). Rather, there are multiple parallel streams of processing going on, as in Dennett's theory of multiple drafts. None of these is 'in' consciousness or

'out' of consciousness; none has a magic extra something that the others lack; none has been rendered explicit or brought into consciousness; but all are phenomenally conscious to the extent that they create some sort of phenomenal world. They arise and fall away but with no one who experiences them. Then, every so often, something different happens. Maybe we want to describe what is going on to ourselves or someone else; or a dramatic event, like a near miss while driving, makes us review our recent experience; or the sudden stop of a ticking clock redirects our attention towards what was or is going on. Then, and only then, is an experiencing self and a briefly unified stream of experiences concocted, making it seem as though we have been conscious all along. At these times, recent events from memory (of the street or the ticking clock, for example) are brought together by paying attention to them, and the appearance of a unified self having unified experiences is created. As soon as attention lapses, the unity falls apart and things carry on as normal. Just as the fridge door is usually closed, so we are usually in a state of parallel multiple drafts. Only when we briefly open the door is the illusion created that the light is always on.



FIGURE 6.5 • Is the light always on inside the fridge?
Is my consciousness always there even when I'm not asking the question?

SUPERUNITY AND DISUNITY


In the penultimate section of this chapter, we will ask what can be learned about unity by studying cases where consciousness seems more or less unified than usual.

SYNAESTHESIA

In synaesthesia, people hear sounds in colour, see shapes in tastes, listen to touches on their skin, or feel tactile sensations on their own skin when seeing someone else being touched. Grapheme–colour synaesthesia is the most common form, and brain imaging has shown that the associative areas at the boundary between the language and visual systems play a key role. In particular, occipital cortex, which is active when we read real words but not non-word strings of letters, offers a possible neural basis for how linguistic stimuli are systematically linked with sensory visual attributes like colour.

Sometimes described as ‘a special case of integrated cross-modal perception’ (Frith and Paulesu, 1997, p. 124), synaesthesia is arguably only a heightened version of what our minds do all the time: combine visual, kinaesthetic, and vestibular signals to track our own bodies in space; attribute olfactory inputs to taste; make consistent multi-sensory associations between high-pitch tones and brighter/lighter colours, or between warmth and affection or light and truth. Maybe the divisions between sensory modalities are not as neat as we tend to assume, implying that we may be wrong to make a mystery out of conscious unity by assuming that that everything is separate to begin with and somehow needs unifying. ‘[T]he existence of synesthesia invites psychologists to reconsider their notions of what “normal” is’ (Ward, 2013, p. 51).

By contrast, there are situations in which the unity of consciousness is lost, whether briefly or lastingly. Some of the most dramatic instances are those of multiple personality and the effects of splitting the brain by cutting the corpus callosum. We will consider multiple personality in Section Six, in connection to concepts of self, and will here focus on the more directly brain-related phenomena of disunity: split brains, amnesia, and neglect.



CONCEPT 6

SYNAESTHESIA

- ‘What a crumbly, yellow voice’, said S. ‘I can’t escape from seeing colors when I hear sounds. What first strikes me is the color of someone’s voice’. S was the famous ‘mnemonist’, or memory man, studied by the great Russian psychologist Aleksandr Luria. S could remember vast tables of numbers and learn poems in languages he did not understand, yet he found communication difficult, and could not hold down a job or forget the pains of his childhood; ‘for S. there was no distinct line, as there is for others of us, separating vision from hearing, or hearing from a sense of touch or taste’ (Luria, 1968, pp. 24–25, 27).

In synaesthesia, events in one sensory modality induce vivid experiences in another. In the most common form, grapheme–colour synaesthesia, written letters or numbers are seen as coloured, but people can hear shapes, see touches, or even have coloured orgasms (Cytowic, 1993). The experiences are vivid and precise and cannot be consciously suppressed, and when tested after many years, most synaesthetes report exactly the same forms or colours induced by the same stimuli (Cytowic and Eagleman, 2009).

Many synaesthetes hide their special abilities, so it is difficult to know how common synaesthesia is. In the 1880s Galton estimated it at 1 in 20, whereas other estimates range from 1 in 200 to 1 in 100,000 (Baron-Cohen and Harrison, 1997). Many people experience days, months, numbers, and the alphabet in a spatial form such as spirals or circles, and this is arguably a weak form of synaesthesia. Synaesthesia runs in families, is more common in left-handers, and is six times more common in women than men. It is associated with artistic ability and good memory but poorer maths and spatial ability.

Synaesthesia has often been dismissed as fantasy, overly concrete use of metaphor, or exaggerated childhood memory, but none of these ideas can explain the phenomena. The induced colours produce cross-modal Stroop interference (Ward, Huckstep, and Tsakanikos, 2006)

and pop-out in complex displays, so that synaesthetes can detect concealed shapes, such as triangles or squares, more easily than controls (Ramachandran and Hubbard, 2001). This shows that synaesthetes are not confabulating or relying on memory for associations.

Cytowic emphasises the connection with emotions and the limbic system, while Ramachandran suggests that shape–colour synaesthesia is caused by a mutation that creates cross-activation between visual areas (especially V4 and V8) and the number area, which lie close together in the fusiform gyrus. Other kinds of synaesthesia may depend on crossover between other neighbouring areas of sensory cortex. Synaesthesia is more common in autism, suggesting that both may result from neural over-connectivity (Baron-Cohen et al., 2013).

SPLIT BRAINS, SPLIT CONSCIOUSNESS?

Epilepsy can be a debilitating disease, at its worst causing almost continuous seizures that make a fulfilling life impossible. For such serious cases, a drastic operation was carried out many times in the 1960s, before less invasive treatments were discovered. To prevent seizures spreading from one side to the other, the two halves of the brain were separated in an operation known as a commissurotomy. In some patients, only the corpus callosum, or part of the corpus callosum, was cut; in others the smaller anterior and hippocampal commissures were also cut. Remarkably, these patients recovered well and seemed to live a relatively normal life. Tests showed that their personality was little changed and their IQ and verbal and problem-solving abilities were hardly affected (Sperry, 1968; Gazzaniga, 1992), but in the early 1960s some clever experiments were designed to test the two hemispheres independently. The

findings about the dramatic consequences of this disconnection, and the work that followed, earned a Nobel prize for pioneering psychobiologist Roger Sperry.

Information from the left visual field goes to the right hemisphere and from the right visual field to the left hemisphere (note, this is not from the left and right eyes). The left half of the body is controlled by the right hemisphere (and vice versa), but information from the right ear goes to the right hemisphere (and vice versa). Knowing this, it is possible to feed information to only one hemisphere, and obtain a response from only one hemisphere. In 1961 neuroscientist Michael Gazzaniga first tested the split-brain patient W. J. using such a procedure. At that time, research on cats and monkeys had shown that the two hemispheres appeared to function almost entirely separately when disconnected, but no one expected this to be true of humans – after all, the patients appeared

to act and speak and think like ordinary unified people. But the research showed that, as in the animals, each half-brain appeared to behave independently.

In a typical experiment, the patient fixated on the centre of a screen which was divided into two. Words or pictures were then flashed to either visual field, thus sending information to only one hemisphere. The patient responded verbally, or




FIGURE 6.6 • Luria (1968) read the following sentence to S.: 'If carbon dioxide is present above a vessel, the greater its pressure, the faster it dissolves in water.' S. was so distracted by the mental images associated with each word that he could not understand this simple rule.

by using one hand or the other to indicate an answer. Suppose that a picture of an object was flashed to the right visual field. Since in most people verbal ability is restricted to the left hemisphere, the patient could then say exactly what it was. But if it was flashed to the left side, he could not. In other words, the left hemisphere, with its ability to control speech, 'knew' the correct answer only when the picture appeared on the right. In anyone with an intact corpus callosum the information would quickly flow across to the other side of the brain, but in these split-brain patients it could not. The interesting finding was that the right hemisphere could communicate in other ways. So if a pile of objects was given, out of sight, to the left hand, that hand could easily retrieve the object seen in the left visual field.

Tasks could even be done simultaneously. For example, asked to say what he had seen, a patient might reply 'bottle', while his left hand was busy retrieving a hammer from a heap of objects – or even retrieving a nail as the closest association. When a dollar sign was flashed to the left and a question mark to the right, the patient drew the dollar sign, but when asked what he had drawn he replied 'a question mark'. As Sperry (1968) put it, one hemisphere does not know what the other is doing; each can remember what it has seen but these memories are inaccessible to the other. This means the left hand could retrieve the same object an hour later, but the person speaking through the verbal left hemisphere would deny ever having seen it.

Sperry thought that these results revealed a doubling of conscious awareness, and even that his patients had two free wills in one cranial vault. 'Each hemisphere seemed to have its own separate and private sensations', he said, and he concluded that 'the minor hemisphere constitutes a second conscious entity that is characteristically human and runs along in parallel with the more dominant stream of consciousness in the major hemisphere' (1968, p. 723). In other words, for Sperry a split-brain patient is essentially two conscious people.

Koch agrees, claiming that 'split-brain patients harbor two conscious minds in their two brain halves', and asking 'How does it *feel* to be the mute hemisphere, permanently encased in one skull in the company of a dominant sibling that does all the talking?' (2004, pp. 294, 293).



ACTIVITY 6.1

Are you a synaesthete?

If you have a large class or other group of people that you can easily test, you can ask people whether they ever experience one sense in response to another, or whether they used to do so as a child. Some people can describe vivid memories of seeing coloured music, or experiencing tastes and smells as having a particular shape, even though they can no longer do so. You may find people who claim extravagant associations and florid experiences. Here are two simple tests that might help detect whether they are making it up or not.

1 Retesting associations. This test needs to be done over two separate sessions, without telling participants that they will be retested. In the first session read out, slowly, a list of numbers in random order (e.g. 9, 5, 7, 2, 8, 1, 0, 3, 4, 6) and a list of letters (e.g. T, H, D, U, C, P, W, A, G, L). Ask your group to visualise each letter or number and write down what colour they associate with it. Some will immediately know, while others may say they are just making up arbitrary associations. Either way they must write down a colour. Collect their answers and keep them. In a second session (say a week or several weeks later), read out the same letters and numbers but in a different order (e.g. 6, 3, 8, 1, 0, 9, 2, 4, 5, 7; P, C, A, L, T, W, U, H, D, G). Give them back their previous answers and ask them to check (or to check a neighbour's) and count how many answers are the same. True synaesthetes will answer almost identically every time they are tested.

2 Pop-out shapes. Tell the group that you will show them a pattern in which a simple shape is hidden. When they see the shape they are to shout out 'Now'. Emphasise that they must NOT say the name of the shape and give the game away, but must just shout 'Now'. As soon as you show the pattern (Figure 6.7), start timing as many of the shouts as you conveniently can. If you have any synaesthetes in the group, they will see the pattern much sooner than everyone else. Even if you have no synaesthetes, these figures can help everyone else to imagine what synaesthesia is like.




FIGURE 6.7 • Schematic representation of displays used by Ramachandran and Hubbard (2001) to test whether synaesthetically induced colours lead to pop-out. (a) A triangle of 2s is embedded in a matrix of 5s. Non-synaesthetes found it hard to find the triangle. Synaesthetes who saw 5s as (say) green and 2s as red found the triangle easily. (b) A single grapheme presented in the periphery is easy to identify but when flanked by others becomes much harder to detect. Synaesthetic colours (like normal colours) can overcome this effect.

At first, Gazzaniga also believed that consciousness had been separated to give a 'double conscious system' (1992, p. 122). But later he began to doubt this conclusion with his discovery of what he called 'the interpreter', located in the left hemisphere. In one test a picture of a chicken claw was flashed to the left hemisphere and a snow scene to the right. From an array of pictures the patient, P. S., then chose a shovel with the left hand and a chicken with the right. When asked why he replied, 'Oh, that's simple. The chicken claw goes with the chicken, and you need a shovel to clean out the chicken shed' (p. 124).

This kind of confabulation was common, especially in experiments with emotions. If an emotionally disturbing scene was shown to the right hemisphere, then the whole body reacted appropriately with, for example, blushing, anxiety, or signs of fear. When asked why, the uninformed left hemisphere always made up some plausible excuse. When the right hemisphere was ordered, for example, to laugh or walk, the whole body would obey. When asked why, the patient might reply that the experimenters were funny, or that he wanted to fetch a Coke. The patients never said things like 'Because I have a split brain and you showed another picture to the other half'.

Don't think that confabulation is confined to neurological patients. Arguably every explanation we give is a confabulation in the sense that even with our two hemispheres connected, we do not know either the neural bases of our actions or all the environmental influences on them. So we have to make sense of them in terms of invented desires, beliefs, opinions, and reasons. As psychologist Steven Pinker puts it,

The spooky part is that we have no reason to think that the baloney-generator in the patient's left hemisphere is behaving any differently from ours as we make sense of the inclinations emanating from the rest of our brains. The conscious mind – the self or soul – is a spin doctor, not the commander in chief.

(2002, p. 43)

'split-brain patients harbor two conscious minds in their two brain halves'

(Koch, 2004, p. 294)

'The same brain may subserve many conscious selves'

(James, 1890, i, p. 401)

And unity is a crucial part of the spin doctor's message.

Where does that leave the non-dominant hemisphere? Is it conscious? Gazzaniga argued that only the left-hemisphere interpreter uses language, organises beliefs, and ascribes actions and intentions to people. So only this hemisphere has what he calls 'high-level consciousness'. Sperry wondered whether the non-dominant hemisphere has 'a true stream of conscious awareness' or is just an 'automaton carried along in a reflex or trance-like state' – what some might call a zombie (1968, p. 731). Further research showed that the two hemispheres have very different abilities: for example, the left has far superior language skills and the right has better face recognition. The right hemisphere has been described as having the linguistic ability of a three-year-old child, or the reasoning capacity of a chimpanzee. Yet we often ascribe consciousness to young children and other animals. In those rare cases where a person's entire left hemisphere is destroyed or removed, we still consider them to be conscious.

Scottish neuroscientist Donald MacKay (1987) was determined to find out whether split-brain patients are really two persons or one, and devised an ingenious test. He taught each hemisphere, separately, to play a 'twenty questions'-type guessing game with him. One person chooses a number from 0 to 9 and the other has to guess what it is by saying 'up', 'down', or 'OK' until the correct answer is reached. Both halves of patient J.W. learnt the game easily. Then they were asked to play against each other, with J.W.'s mouth (controlled by his left hemisphere) making the guesses, and his left hand (controlled by his right hemisphere) pointing to cards saying 'go up', 'go down', or 'OK'. With this game, it proved possible for the two half-brains to play against each other, and even to cooperate and pay each other winnings in tokens, but MacKay concluded that there was still no evidence of two separate persons or of true 'duality of will'.

How, he asked, could anything play a game of twenty questions without being conscious? He noted all the intelligent actions we can carry out unconsciously, and the artificial systems that can play games, and came to the following conclusion. To understand human behaviour we must distinguish between the executive and supervisory levels of brain function. The executive level can (unconsciously) control goal-directed activities and evaluate them in terms of current criteria and priorities, but only the self-supervisory system can determine and update those priorities. We are conscious only of those features of our world that engage this self-supervisory system.

With this theory, MacKay provides his own answers to some of our recurring questions about consciousness. Question: what makes some things conscious and others not? Answer: whether they engage the self-supervisory system or not (though he admits that how the activity of this system gives rise to conscious experience remains totally mysterious). Question: what makes each of us a psychological unity? Answer: that we have only one self-supervisory system to determine our overall priorities. As for the split-brain patient: he has only one self-supervisory system and therefore is still only one conscious person.

So who was right? Do split-brain patients have one consciousness or two? At first sight, it seems as though there must be an answer to this simple question, but perhaps we need to think again.

You may have noticed that in the paragraphs above we sometimes described one hemisphere or the other as knowing, seeing, or even being conscious. This kind of language is hard to avoid when confronted with the strange findings we have described, but the philosophers Maxwell Bennett and Peter Hacker (2003) accuse Sperry, Gazzaniga, and other neuroscientists of making conceptual errors and causing 'profound confusion' by talking about a half brain as though it were a person.

It should be obvious that the hemispheres of the brain can neither see nor hear. They cannot speak or write, let alone interpret anything or make

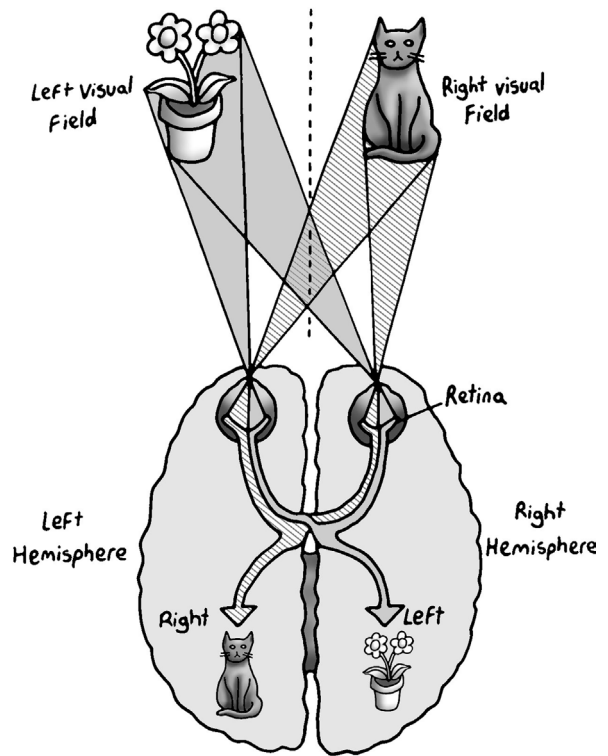


FIGURE 6.8 • The human visual system is organised as shown. Information from the left visual field of both eyes (in this case the flowers) goes to the right hemisphere, and information from the right visual field of both eyes (in this case the cat) goes to the left hemisphere. Note that by this partial crossing-over of fibres in the optic chiasm, the effect is that the two sides of the brain deal with opposite sides of the world, not with opposite eyes.

SECTION TWO: THE BRAIN



FIGURE 6.9 • The split-brain patient P.S. was shown a snow scene to the right hemisphere and a chicken claw to the left hemisphere, and asked to choose from an array of pictures. He chose the shovel with his left hand and the chicken with his right (Gazzaniga, 1992, p. 127).

'the life of a second rudimentary self lasts a few minutes at most'

(Dennett, 1991, p. 425)

inferences from information. They cannot be said to be either aware or unaware of anything.

(Hacker and Bennett, 2003, p. 391)

Only whole human beings can be said to be conscious. In their view, a split-brain patient is deprived of the capacity to carry out normally coordinated functions, not split into two people. This is just one example of what Bennett and Hacker call the 'mereological fallacy': the widespread and almost unquestioned tendency of neuroscientists to say that brains or parts of brains can see, hear, think, make decisions, or experience things when all these abilities are functions of whole human beings, not bits of brain.

Along similar lines, philosopher Michael Tye (2003) argues that split-brain patients are persons whose phenomenal consciousness is briefly split during the experiments but is otherwise unified, whereas Tim Bayne (2005) responds that although a *person* might have two streams of consciousness, it doesn't make sense to think of *subjects of experience* as doing so.

We have met several theories that assume a difference between 'conscious' and 'unconscious' processes and try to explain it, and Zeki specifically links his micro-consciousnesses to specific nodes in the brain.



FIGURE 6.10 • When the silent right hemisphere is given a command, it carries it out. At the same time, the left doesn't really know why it does so, but it makes up a theory quickly (reprinted from Gazzaniga and LeDoux, 1978, in Gazzaniga, 1992, p. 128).

Perhaps the best use we can make of the mereological fallacy is to notice whenever people attribute consciousness to parts of a human being and ask ourselves whether this is really what they intended.

AMNESIA

‘What year is this, Mr G.?’ asked the neurologist and author Oliver Sacks (1985, p. 25).

‘Forty-five, man,’ his patient replied. ‘What do you mean? We’ve won the war [. . .]. There are great times ahead.’

‘And you, Jimmie, how old would you be?’

‘Why, I guess I’m nineteen, Doc.’

Sacks then had an impulse for which he never forgave himself. He took a mirror and showed the 49-year-old grey-haired man his own face. Jimmie G. became frantic, gripping the sides of his chair and demanding to know what was going on. Sacks led him quietly to the window where he saw some kids playing baseball outside. Jimmie G. started to smile and Sacks stole away. When he returned a few minutes later, Jimmie greeted him as a complete stranger.

Jimmie G. had Korsakoff’s syndrome, and nothing can be done to restore memory in such cases. Jimmie first lost his ability to form new long-term memories (anterograde amnesia), and then began to lose his long-term memory for past events (retrograde amnesia). Jimmie’s amnesia was ‘a pit into which everything, every experience, every event, would fathomlessly drop, a bottomless memory-hole that would engulf the whole world’ (Sacks, 1985, p. 35).

This is not, however, a complete loss of all memory. Classical conditioning remains unimpaired, so that patients with Korsakoff’s easily learn to blink to a sound if it is paired with a puff to the eye; to associate certain smells with lunchtime; or to respond to a given visitor with pleasure, even if they claim never to have seen that person before. Procedural learning also remains intact. Not only do people with amnesia often retain such skills as driving a car or typing, they can learn new ones. They might, for example, learn to mirror-read but be unable to remember the words they read, and even deny ever having learned the skill; or they might improve at playing a computer game without remembering they have played it before. They



ACTIVITY 6.2

Split-brain twins

Ask for two volunteers: one to play the role of a disconnected left hemisphere (LH) and the other to play the right (RH). Ask them to sit close together on a bench or table. You might like to put a sticker on each, labelling them as LH and RH. To reduce confusion we’ll assume for this explanation that LH is female and RH is male.

LH sits on her left hand; her right hand is free to move. RH sits on his right hand; his left hand is free to move. Their two free arms now approximate to those of a normal person. RH cannot speak (although we will assume that he can understand simple verbal instructions). You might like to tape his mouth over, making sure the tape will not hurt when removed.



FIGURE 6.11 • Experiment 1: RH puts his free left hand in the bag and feels the object. When you ask what he feels, only LH can speak.

Now you can try any of the split-brain experiments described in this chapter. Here are just two examples.

1 You will need a large carrier bag or a pillow case containing several small objects (e.g. pen, shoe, book, bottle).

Out of sight of LH, show RH a drawing of one of your objects. Ask him ‘What can you see?’ Only LH can speak and she did not see the drawing. Press her to

answer (if RH tries to give her non-verbal clues this only adds to the fun). Now give RH the bag containing the various objects, ask him to put in his left hand and select the correct object. He should do this easily.

2 Recreate MacKay's experiment.

Ask RH to think of a number between 0 and 9 (or, if you want better control, prepare numbered cards and show one to RH out of sight of LH). LH now has to guess what the number is. For each guess RH points 'up' or 'down' or nods for the correct answer. You might like to try inventing a method for playing the game the other way round.

The twins should be able to play this game successfully. Does it show that there are two conscious selves involved? Does this game help us to understand what it is like to have a split brain?

'I have just become conscious for the first time'

(C. W., amnesic patient)

also show evidence of priming: getting quicker at recognising fragmented pictures and completing words if they have been seen before. For this reason, amnesic syndrome has sometimes been described as a dissociation between performance and consciousness (Farthing, 1992; Young, 1996).

Are people with amnesia conscious? Surely the answer is yes. They are awake, responsive, able to converse, laugh, and show emotion. But without the capacity to create new memories, they have lost the interaction between current and stored information that, according to Oxford psychologist Larry Weiskrantz (1997), makes possible the 'commentary' that underlies and unifies conscious experience. Some amnesiacs repeatedly exclaim 'I have just woken up!' or 'I have just become conscious for the first time!'

C. W. was a professional musician struck with dense amnesia by herpes simplex encephalitis (Wilson and Wearing, 1995). Although he could still sight-read and improvise music, and even conduct his choir, his episodic memory was almost completely destroyed. He kept a diary of what was happening to him and there he recorded, hundreds of times, over a period of nine years, that he was now fully conscious, as if he had just woken from a long illness. He was conscious all right, but trapped in an ephemeral present, unconnected with the past.

Asking people with amnesia about such matters is difficult. As Sacks puts it,

If a man has lost a leg or an eye, he knows he has lost a leg or an eye; but if he has lost a self – himself – he cannot know it, because he is no longer there to know it.

(1985, pp. 35–36)

Amnesics create no memory of a continuous self who lives their life, or, as some would say, no illusion of a continuous self who lives their life.

Amnesia will come to many of us, and to our parents and loved ones, in the form of Alzheimer's disease or senile dementia. In this form, it is less specific than the cases described here, and comes on gradually. For some time, the person may have enough memory to realise their predicament, which makes it all the harder. In 2014, Bruce Francis, Emeritus Professor of Electrical Engineering at the University of Toronto, gave a prize lecture on 'The Robot Rendezvous Problem', and his first PowerPoint slide read: 'I have Parkinson's disease. To help me

deliver the lecture as smoothly as possible, I've written text on the slides and will read it. You should read it along with me (not aloud). Let's practice by an example.' He then proceeded to give a masterclass in how to present complex ideas from the ground up.

Memory loss can be frightening. Yet, as the Russian psychologist Alexander Luria pointed out to Sacks, 'a man does not consist of memory alone. He has feeling, will, sensibilities, moral being – matters of which neuropsychology cannot speak' (Sacks, 1985, p. 32). Memory is just one of the kinds of glue we use to give unity to our consciousness.




FIGURE 6.12 • Korsakoff's syndrome occurs most often in chronic heavy drinkers. It is caused by thiamine (vitamin B1) deficiency, which damages the thalamus and hypothalamus, as shown, and is exacerbated by the neurotoxic effects of alcohol. This results in anterograde amnesia (the inability to create new memories). Other brain areas including the cerebellum may also be affected.

NEGLECT

Some people who have a stroke causing damage to the right side of the brain lose the left-hand side of their world. In the phenomenon of hemifield neglect, or unilateral neglect, patients seem not to realise that the left-hand side of the world even exists (Bisiach, 1992). After a stroke to the right hemisphere, one woman applies make-up only to the right side of her face, and eats only from the right side of her plate. A man shaves only the right side of his beard, and sees only the right side of a photograph.

Many tests reveal the peculiarities of this condition. When asked to copy a drawing of a flower, some patients accurately copy the right half, while others squash all the petals onto the right side. When asked to draw a clock face, some leave out the left half, while others squash all the numbers onto the right. And when asked to bisect a horizontal line, they typically mark it far to the right of the mid-point. However, it is not as though they have entirely lost half their vision: visual responsiveness remains in the neglected areas, and stimuli that are neglected can prime later responses. Instead, they have lost something much more fundamental.

Italian neurologist Edoardo Bisiach asked his neglect patients to imagine Milan's beautiful cathedral square. First he asked them to imagine standing at one side, facing the fantastic Duomo with its pinnacles and magnificent façade, and to describe what they saw. They knew the Piazza well and described the buildings that would lie to their right when standing in that position, leaving out all those on the left. But they had not forgotten the existence of those on the left. When asked to imagine standing on the other side, facing the other way, they described all the buildings they had previously left out (Bisiach and Luzzatti, 1978). Although they have thorough knowledge of the buildings on both sides, when they imagine the square, the left side simply doesn't exist.

Hemifield neglect can partly be explained as a deficit of attention, in that patients simply do not attend, or have their attention drawn, to the left-hand side of their world, and to some extent they can be helped by training them to keep turning from side to side. Yet clearly the unattended side is not completely blanked out. For example, emotional stimuli shown in the neglected field can influence attention. In one experiment, patients were shown two pictures of a house, identical except that one had flames pouring from a window on the left-hand side. While insisting that the houses were identical, patients still said they would prefer to live in the one that was not on fire (Marshall and Halligan, 1988). Although subsequent studies have shown rather different results for the house test, the conclusion remains that stimuli that are not consciously seen can still affect behaviour.

Weiskrantz describes it this way: 'The subject may not "know" it, but some part of the brain does' (Weiskrantz, 1997, p. 26). But perhaps this implies a unitary, superordinate 'subject' who watches the workings of the lower mechanisms. According to Bisiach (1988), there is no such entity, for the task of monitoring inner activity is distributed throughout the brain. When lower-level processors are damaged, higher ones may notice, but when the higher ones are gone there is nothing to notice the lack.

He believed that 'some of the questions set by commissurotomy, blindsight, unilateral neglect of space, etc. will remain forever unanswerable: without direct

'we will never know what it is like to be a patient affected by unilateral neglect'

(Bisiach, 1988, p. 117)




FIGURE 6.13 • Two drawings made by a patient suffering from unilateral neglect, studied by Marshall and Halligan (1988). The drawing on the left was done in the acute phase after the stroke, while the one on the right was drawn in the chronic phase several months later.

experience we will never know what it is like to be a patient affected by unilateral neglect' (1988, p. 117). But perhaps we already do. Our eyes and ears detect only a small range of wavelengths; we have no electrical sense like some fish nor infra-red detectors like some snakes. From the richness of the world out there, our senses select what they are evolved to select, and the rest we do not miss and cannot even imagine. This is why Metzinger (2009) describes us as living in a tunnel. In this sense, we all live our lives in a profound state of neglect.

Other examples that challenge the unity of consciousness include out-of-body and near-death experiences, in which consciousness seems to split from the physical body (Chapter 15); and mediumship, trances, and hypnosis (Chapter 13), in which consciousness can seem to be divided. Although many people assume that consciousness is necessarily unified most of the time, there are plenty of reasons for doubt.

It always seems to me that our ordinary consciousness inhabits the tip of a pyramid whose base extends so widely within us (and to a certain extent under us) that the further we think we're able to let ourselves sink into it, the more wholly we seem encompassed by the timeless and spaceless givenness of earthly, in the broadest sense worldly, existence. Since my youth I have always suspected [...] that in a deeper section of this pyramid of consciousness, simple existence could become an event – that unbreakable presentness and simultaneity of everything which at the 'normal' pinnacle of self-consciousness can only be experienced as 'sequence'.

(Rainer Maria Rilke (1980), letter to Nora Purtscher-Wydenbruck, 11 August 1924, our translation)




FIGURE 6.14 • Figure used to investigate covert processing in a patient with hemineglect of the left half of the visual space. The two figures looked identical to the patient because only their right halves were reported as seen. Nevertheless, when required to indicate which house she would prefer to live in, she chose the bottom one, although she said she was guessing (Marshall and Halligan, 1988).

UNITY IN TIME

We tend to feel not only that everything comes together 'in' our conscious experience, but that it happens *now*, and we may think that the *now* of our consciousness is the same as the *now* of the world we're perceiving; we don't need to ask ourselves 'when is now?' or what it means to say that there is a time at which conscious experiences happen. But some experiments challenge these simple intuitions.

Clock time and experienced time are not the same thing. While events in the external world can be timed with clocks, as can events happening inside the brain (such as neurons firing), perceived time is not like this. We can only find out about it by asking a person to report in some way. This is notoriously difficult, as we will find out in Chapter 9, when we look at Libet's famous experiments on apparent time lags in both conscious sensation and willed action.

Wilhelm Wundt did early experiments on time and sensory consciousness towards the end of the nineteenth century. He asked people to judge the relative timings of visual and auditory stimuli and found many examples of what he called 'subjective time displacement', in which people made mistakes about which event occurred first. In the 1920s, F. W. Fröhlich observed that if a moving object suddenly appears, its perceived initial location is misplaced in the direction of motion.

Modern illusions include the flash-drag, flash-jump, and flash-lag effects. In the last of these, one object moves continuously and another is flashed just as it is aligned with the moving one. The flash then appears to lag behind the moving object. One proposed explanation is that the visual system predicts where the moving stimulus is going so as to allow for processing delays. Another is that processing is done 'online' but moving objects are processed more quickly than static ones. Eagleman and Sejnowski (2000) argue that neither of these explains the effect, and propose that visual awareness is not predictive but postdictive, so that events which happen shortly after the flash (within about 80 milliseconds) affect what is perceived.

These and many other experiments make the point that we do not always experience things, or report their occurrence, in the order in which they actually happen in the world. From this we might be tempted to imagine something like this: there are two worlds—a physical world in which events really happen in one order, and an inner experienced world of consciousness in which they happen in another order. This dualist view is tempting, but illusions like this do not necessarily imply duality, and some further phenomena will help to show just how problematic it is.

If two lights in different positions are flashed quickly one after the other, there appears to be one light moving, rather than two separate lights flashing. This is the well-known phi phenomenon. In 'colour phi', the lights are different colours, say red and green. In this case, something very odd happens. Observers often report that the light not only moved but also changed from red to green as it did so. How can this be? The light seems to start changing colour before the second flash, but how could the person know that a green light was coming?

A similar problem occurs with the 'cutaneous rabbit' (Geldard and Sherrick, 1972; Dennett, 1991) (see Activity 6.3). If a person's arm is tapped, say five times at the wrist, twice near the elbow, and then three times on the upper arm, they report not a series of separate taps coming in groups, but a continuous series moving upwards – rather as though a little creature were running up their arm. Once again, we might ask how taps two to four came to be experienced as moving up the arm when the next tap in the series had not happened yet. How did the person know where the next tap was going to be?

This certainly seems mysterious, so what is going on? We might, perhaps, think that colour phi works like this: first the person consciously experienced a stationary red light, then when the green light flashed this experience was wiped out and replaced with the new experience of the light changing to green. Alternatively, we might suppose that the person never did consciously experience the

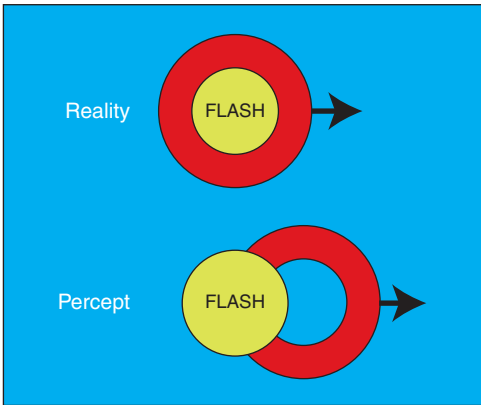


FIGURE 6.15 • The flash-lag illusion. When a flash and a moving object are shown in the same position they appear to be displaced from each other. In this version the flash appears to lag behind the moving ring.

stationary red light because consciousness was delayed until all the relevant information was in, and only then was it allowed 'into consciousness'.

Dennett investigates many such phenomena and asks how we might distinguish between these two views. Surely one must be right and the other wrong? Surely we must be able to say, at any point in time, what was actually in that person's stream of consciousness, mustn't we? No, says Dennett, because there is no way, in principle, of distinguishing these two interpretations. He illustrates this by comparing two fanciful interpretations, the Orwellian versus the Stalinesque revision (Concept 6.2). Ultimately, 'This is a difference that makes no difference' (Dennett, 1991, p. 125). If we think one or other must be true, we are still locked in the Cartesian Theatre.

How then can we understand these odd phenomena? When things seem mysterious, it is often because we are starting with false assumptions. Perhaps we need to look again at the very natural assumption that when we are conscious of something there is a time at which that conscious experience happens.

It may seem odd to question this, but the value of these oddities may lie precisely in forcing us to do so. The problem does not lie with timing neural events in the

brain, which can, in principle, be done. Nor does it lie with the judgements we make about the order in which things happen. It begins when we ask 'But when does the *experience itself* happen?' Is it when the light flashes? Obviously not, because the light hasn't even reached the eye yet. Is it when neural activity reaches the lateral geniculate? Or the superior colliculus? Or V1, or V4? If so, which and why, and if not, then what? Is it when activity reaches a special consciousness centre in the brain (or in the mind)? Or when it activates some particular cells? Or when a complicated consciousness-inducing process is carried out?

Almost all the theories we have encountered so far give answers to these kinds of question. For example, in GWT things become conscious when they enter the global workspace and are broadcast, for Zeki consciousness happens when activity in brain cells becomes explicit, and for Crick the 'awareness neurons' must be active. But any theory of this kind has to explain how subjective experiences arise from this particular neuron, or this particular kind of neural activity, at this particular time.

So perhaps we need to drop yet another intuition – the idea that there must be a time at which conscious experiences happen. Perhaps we in fact create temporal unity retrospectively, in a kind of very personal storytelling. Our stories include the order in which things have happened, but only as a reasonable way to make sense of events, not because any 'actual conscious experiences' also happened in that order.



ACTIVITY 6.3

The cutaneous rabbit

The cutaneous rabbit is easy to demonstrate and a good talking point. You will need a very sharp pencil or a not-too-dangerous knife point – something with a tiny contact point but not sharp enough to hurt. Practise the tapping in advance until you can deliver the taps with equal force and at equal intervals.

Ideally use a volunteer who has not read about the phenomenon. Ask the volunteer to hold out one bare arm horizontally and to look in the opposite direction. Take your pointed object and, at a steady pace, tap five times at the wrist, three times near the elbow, and twice on the upper arm, all at equal intervals. Now ask what it felt like.

If you got the tapping right, it will feel as though light taps ran quickly up the arm, like a little animal. This suggests the following questions. Why does the illusion occur? How does the brain know where to put the second, third, and fourth taps when the tap on the elbow has not yet occurred? *When* was the volunteer conscious of the third tap? Does Libet's evidence help us understand the illusion? What would Orwellian and Stalinesque interpretations be? Can you think of a way of avoiding both?



FIGURE 6.16 • If the disc is flashed briefly (e.g. 30 msec) and immediately followed by the ring, superimposed upon it, the participant reports seeing only the ring. One interpretation is that the ring prevents the experience of the disc from reaching consciousness, as though consciousness is delayed and then changed if necessary (Stalinesque). An alternative is that the disc is consciously experienced but memory for the experience is wiped out by the ring (Orwellian). How can we tell which is right? We cannot, says Dennett. This is a difference that makes no difference.

‘unless there were a Cartesian theater, there could not be a fact of the matter distinguishing Orwellian from Stalinesque content revisions’

(Dennett, 1991, p. 440)



CONCEPT 6.2

ORWELLIAN AND STALINESQUE REVISIONS

Is there a precise moment at which something ‘becomes conscious’ or ‘comes into consciousness’? In *Consciousness Explained*, Dennett says no. Take the simple example of backwards masking. A small solid disc is flashed first, followed quickly by a ring. If the timings and intensities are just right, the second stimulus masks the first, and observers say they saw only the ring.

What is happening in consciousness? If you believe in a time at which a visual experience ‘becomes conscious’ or comes ‘into consciousness’, then you have two explanations to choose from. Dennett named the first after George Orwell’s novel *Nineteen Eighty-Four*, in which the Ministry of Truth rewrote history to prevent people knowing what had really happened. According to this Orwellian explanation, the person really saw, and was conscious of, the disc,

but then the ring came along and wiped out the memory of having seen it. So only the ring was reported. The alternative he named after Stalin’s notorious show trials, in which people testified to things that never actually happened. On this explanation, the experience of the disc is somehow delayed on its way up to consciousness so that before it gets there the ring can come along and prevent the disc from ever arriving.

The difference hinges on the question – did the disc become conscious and get forgotten, or did it never reach consciousness in the first place? Do you think that there must be an answer to this question?

Dennett (1991, pp. 115–126) argues that there is no way, even in principle, that we could find out. So the question is meaningless. He analyses the ways in which Orwellian and Stalinesque explanations have been used and shows that they always end in an impasse. The problem, he says, is a false assumption. We wrongly assume that there is not only a real time at which things happen in the brain, but also a time at which they ‘enter consciousness’ or ‘become conscious’. If we drop this assumption (difficult as it is to do so), the problem disappears. According to his multiple drafts model, different streams of activity may be probed to elicit various responses, but none is ever either ‘in’ or ‘out’ of consciousness. So the problem does not arise.

READING



Engel, A. K. (2003). Temporal binding and the neural correlates of consciousness. In A. Cleeremans (Ed.), *The unity of consciousness: Binding, integration and dissociation* (pp. 132–152). New York: Oxford University Press.

Sets out the theory and evidence for the relations between binding and consciousness.

Sacks, O. *The man who mistook his wife for a hat* (1985, London: Duckworth) or *An anthropologist on Mars: Seven paradoxical tales* (1995, London: Picador).

Read any chapter from either of these books, which offer a neurologist's accounts of unusual psychiatric cases and states, including: memory loss, body-image disturbance, phantom limbs, Tourette's syndrome, autism, colour blindness, and musical prodigies. Students can report on what they think a chapter's implications are for consciousness.

Tononi, G. (2015). Integrated information theory. *Scholarpedia*, 10(1), 464. www.scholarpedia.org/article/Integrated_information_theory


An accessible account of the qualities of consciousness the theory tries to explain, and how it does so.

Ward, J. (2013). Synesthesia. *Annual Review of Psychology*, 64, 49–75.

Outlines current thinking on synaesthesia's characteristics and mechanisms, and its relevance to other aspects of mind, including consciousness.

Zeki, S. (2007). A theory of micro-consciousness. In M. Velmans and S. Schneider (Eds), *The Blackwell companion to consciousness* (pp. 580–588). Oxford: Blackwell.


Unfolds Zeki's theory from the starting point of the visual brain's disunity.



Body and world

SECTION

Three



Attention

CHAPTER SEVEN

'Every one knows what attention is', said William James in 1890.

It is the taking possession by the mind, in clear and vivid form, of one out of what seem several simultaneously possible objects or trains of thought. Focalization, concentration, of consciousness are of its essence. It implies withdrawal from some things in order to deal effectively with others.

(1890, i, pp. 403–404)

'Every one knows what attention is'

(James, 1890, i, p. 403)

'No one knows what attention is', said the psychologist Harold Pashler in 1998 (p. 1). 'There is no such thing as attention', said the psychologist Britt Anderson in 2011. The things that seem obvious about attention and those that get us hopelessly confused are most apparent of all when we ask how attention relates to consciousness.

The very familiarity of the concept of attention can make it hard to think about clearly, but perhaps we should start with how it feels. The metaphor of the 'spotlight of attention' comes easily to mind because paying attention feels this way – like directing a light on some things and not others. Perhaps it feels as though attention makes things brighter, more prominent, or more focused.

These notions have a long history. Writing a little before James, the Scottish metaphysician Sir William Hamilton wrote that 'Attention is consciousness and something more; [. . .] it is consciousness concentrated' (Hamilton, 1895, p. 941).

'No one knows what attention is'

(Pashler, 1998, p. 1)

*'attention is
consciousness and
something more; [. . .]
it is consciousness
concentrated'*

(Hamilton, 1895, p. 941)

James quotes Gustav Fechner suggesting that someone who focuses attention on something does not see its colour as brighter or its sound as louder, but 'feels the increase [in intensity] as that of his own conscious activity turned upon the thing' (James, 1890, i, p. 462). This idea was later adapted by phenomenologists, who explore the structures of consciousness from a first-person perspective. They stressed how attention moulds the structure of consciousness, for example into foreground and background or centre and periphery. This idea survives in recent accounts of how attention provides a kind of 'experiential highlighting' that allows us to track, inspect, and act with respect to another person or object (Campbell, 2002).

The metaphor of the spotlight has found a place in many scientific theories of mind, including Francis Crick's 'astonishing hypothesis' (1994) and global workspace theory (e.g. Baars, 1997a, 1997b). Others have elaborated on the intuitive metaphor, giving us variants like the zoom-lens model (Eriksen and St James, 1986) and the blinking spotlight (VanRullen, Carlson, and Cavanagh, 2007) or doughnut-shaped spotlight (Müller and Hübner, 2002).

These metaphors should not be taken too literally, and have often been criticised, for example on the grounds that attention simply improves access to, or decision-making about, what is already represented 'in' visual consciousness. There has been much discussion of whether or not attention increases brightness contrast, or merely improves the accuracy of our perceptions by making us process things more deeply (Prinzmetal et al., 2008). Ned Block has also argued that such changes in experience should be thought of as changing not the content of the experience but the nature of the 'mental paint' we apply when paying attention (Block, 2010).

Yet experiments have found that the metaphor of 'lighting up' has a literal basis: a real attentional 'spotlighting' effect in visual perception. Participants kept their eyes fixated on the fovea (where spatial resolution is highest) and were shown textures in the periphery (where it is much lower).

When they attended to the textures (still not moving their eyes), they could more easily distinguish them (Yeshurun and Carrasco, 1998). It was as though their spatial resolution had improved. Crucially, in tasks where enhanced resolution actually makes the task harder, this effect was found for focused attention, too: participants' performance got worse. Later experiments found that the same effect for brightness, contrast, and colour saturation, but not for differences in hue (Fuller and Carrasco, 2006). It seems that, as in James's notion of focalisation and concentration, attention actually increases the spatial resolution of what we see. It may also change visual and other sensory experience in different ways depending on

context, so it seems that attention can qualitatively shape the kinds of conscious experiences we have – even if, as James also pointed out, we know how to adjust for these effects so we are not misled into thinking the light actually just got brighter.



FIGURE 7.1 • Attention feels like a searchlight in the attic, lighting up now the objects right in front of us, and then some long-forgotten memory from the darkest corner of our mind.

DIRECTING ATTENTION

The image of the spotlight of attention is tempting, but perhaps more careful attention to one's own experience might provide different metaphors. This is one way for 'first-person practice' to feed into the science of consciousness and one reason why we ask you to devote time and energy to the 'Practices' suggested in each chapter: we cannot hope to understand consciousness in general unless we are familiar with our own personal version of it. And as the idea of 'paying careful attention to experience' implies, attention itself is at the heart of all such practice. We will begin with a basic element of our everyday experience of attention, the directing of attention, and ask what basic facts we can establish about it.



PRACTICE 7.1

DID I DIRECT MY ATTENTION OR WAS MY ATTENTION GRABBED?

As many times as you can, every day, ask yourself '***Did I direct my attention or was my attention grabbed?***'

You might begin by asking the question every time whenever you realise that you are attending to something and don't know why. With practice you may find that you can do it for much of the time. This way you can learn to watch the process and come to appreciate how and when your attention shifts. Keep a record of the effects this has on your awareness.

Imagine you are sitting in a lecture and the door opens. You turn round to see who it is. What has happened? If someone asked you, you might say, 'I heard the door open and so I turned round to see who it was.' The causal sequence seems to be: 1) consciously hear sound; 2) turn round to look. It *feels* as though our conscious perception of the noise, possibly followed by a conscious decision to pay attention, is what *caused* us to turn around and pay attention. Is this right? Does conscious perception or conscious will *cause* attention to be directed to a specific place? If it does not always do so, can it ever do so?

First, it seems clear that conscious effort and perception are not always required. Attention can be involuntarily grabbed or intentionally directed, and these processes depend largely on different systems in the brain. Attention is drawn involuntarily when we react quickly to something like a loud noise, or our name being called, or an email notification on our phone, and only realise afterwards that we have done so. Such involuntary attention depends on the ventral attention system, which includes alerting and vigilance systems and is found mainly in the right hemisphere in frontal, parietal, and temporal areas. By contrast, when we deliberately pay attention to someone speaking, or try to ignore an annoying noise to concentrate on reading our book, this uses the dorsal attention system.

• SECTION THREE: BODY AND WORLD

This is found bilaterally in frontal and parietal areas and mediates purposeful, voluntary, or high-level attention. It includes response systems in the prefrontal cortex and anterior cingulate gyrus, as well as orientating systems in the posterior intraparietal sulcus and frontal eye fields. (Note that these systems, in cingulate cortex and frontal areas, are distinct from the dorsal and ventral streams in the visual system, which originate in primary visual cortex in the occipital lobe at the back of the brain and run forwards to the parietal lobe and down into the temporal lobe respectively.)

Flexible control of attention needs dynamic collaboration of both the ventral and dorsal attention systems to balance 'top-down' goals with 'bottom-up' sensory inputs (Vossel, Geng, and Fink, 2014), perhaps within the wider context of feedback-driven probabilistic inferencing about the world (Ransom, Fazelpour, and Mole, 2017). And the categories themselves also aren't simple: within 'top-down' attentional selection, current goals may lose out to reward associations based on past selection history, for example (Awh, Belopolsky, and Theeuwes, 2012). Neat opposites are easy and often dangerous tools of thought (Anderson, 2011, provides a list of twelve pairs for attention alone). But fMRI studies do show that the basic functional organisation of the two systems can be seen even when there are no external demands (Fox et al., 2006).

An important example of involuntary attention is the bottom-up control of eye movements. Our eyes constantly jump around from one fixation point to another. These movements are called saccades and happen several times a second, whether we are aware of them or not. We can also control saccadic eye movements voluntarily and this involves primarily cells in the superior colliculus. If a bright, salient, or moving object is detected in the periphery, the eyes quickly turn to bring that part of the visual world onto the fovea. This must be done very fast to be useful to a moving, acting animal and, not surprisingly, much of the control is coordinated by parts of the dorsal visual stream, in particular the posterior parietal cortex.

In 'smooth pursuit', the eyes can track a moving object, keeping its image on roughly the same part of the fovea. This kind of eye movement is hard to make without an actual moving target and is affected by drug use and by conditions such as schizophrenia, autism, and post-traumatic stress. Oddly, it can continue without conscious awareness, as was shown in experiments with a man who was cortically blind; that is, he was blind because of damage to his visual cortex, while his eyes and other parts of the visual system remained intact. He could not consciously see movement at all, and when surrounded with a large moving stripe display he denied having any visual experience of motion. Yet his eyes behaved relatively normally in tracking the moving stripes, making slow pursuit movements followed by rapid flicks to catch up (Milner and Goodale, 1995, p. 84). This showed that although movement may be necessary for accurate pursuit, awareness of the movement is not.

Nonetheless, even with actions as apparently involuntary as smooth pursuit, the story is complicated. If you know which way a target will move, or know when the motion will begin, you can initiate smooth pursuit before any movement happens. You can keep it going if the moving target is temporarily hidden by another object, and if you move your hand in the dark, the proprioceptive motion signal replaces the visual signal. So even with something as apparently simple as

perceptual pursuit of a moving object, the relationship between consciousness and attention gets rapidly more complex.

And, of course, moving the eyes is not the whole story: the head and body move as well, so there must be mechanisms for coordinating all these movements. For example, information from the motor output for body and eye movements can be used to maintain a stable relationship to the world, even while the body, head, and eyes are all moving. Some control systems appear to be based on retinocentric coordinates – keeping objects stable on the retina – while others use craniocentric coordinates: keeping the world stable with respect to the head. Although we can voluntarily control body and head movements, and some kinds of eye movements, most of the time these complex control systems operate very fast and unconsciously. These are just some of the mechanisms that would be involved when you turned round to see who was coming in through the door.

Another form of involuntary visual attention occurs in perceptual ‘pop-out’. Imagine you are asked to search for a particular stimulus which is displayed amongst a lot of slightly different stimuli, say an upside-down L amongst a lot of upright Ls. For many such displays, there is no alternative but a serial search, looking at each item in turn to identify it. In other cases, the difference is so obvious to the visual system that the target just pops out, for example when the target L is horizontal or is a different colour. In these cases the search seems to be parallel and does not take longer if the total number of items increases. An obvious item like this can also act as a distractor, slowing down the search for other items – another example of how attention can be grabbed involuntarily.

Directing the eyes towards a particular object is not, however, equivalent to paying attention to it. This is true for several reasons. First of all, it is perfectly possible to be blind to something we are looking right at, just because we are not attending to it. In Chapter 3 we learned about the discovery of inattentional blindness, beginning with Arien Mack and Irvin Rock’s work in the late 1990s, and expanding to investigate the role of characteristics like familiarity, expectation, and different kinds of salience in determining whether or not inattentional blindness is experienced.

Other kinds of blindness are an integral part of paying attention. Attention always has costs as well as benefits. Not only does directing attention to one thing mean you have to neglect another, but there may be a short ‘attentional blink’ afterwards. This has been shown in experiments where, for example, a series of letters are rapidly flashed and participants asked to look for a given target letter. If they successfully detect one, then they are less likely to detect another shown within 200–500 msec after the first, as though their capacity to attend ‘blinked’ for a moment, even though they were looking right at the relevant stimulus.

It is also possible to attend visually to two different locations at once. In one fMRI study, participants were asked to fixate a central point while selectively attending to two different targets on either side; they were presented with a task-irrelevant sequence of digits at the central fixation point, and had to identify matching digits from rapid sequences of letters and digits in the left and right locations. Activation in the retinotopic maps in primary visual cortex was found corresponding to both spots but not to the central stimulus in between, suggesting not just one spotlight of attention but ‘multiple spotlights of attentional selection’ (McMains




FIGURE 7.2 • Search for the two odd ones out in each picture. In the top one you will probably have to do a serial search, looking at each L in turn. In the bottom picture the horizontal Ls just pop out.

DID I DIRECT MY ATTENTION OR WAS IT GRABBED?

‘There is no conscious perception without attention’

(Mack and Rock, 1998, p. 14)

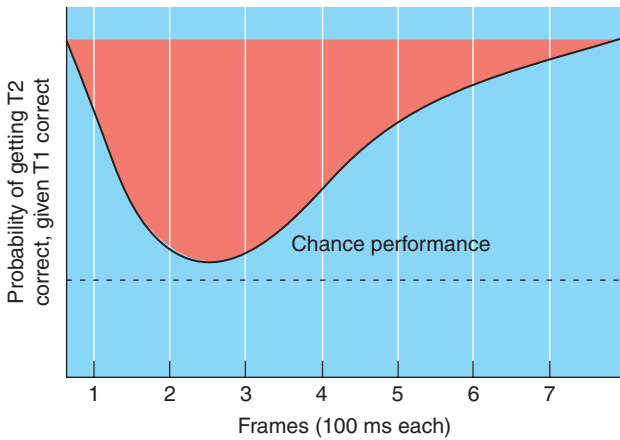


FIGURE 7.3 • An example of a prototypical procedure used to measure the attentional blink (AB). (a) Depiction of the experimental design. The targets are numbers and distractors are letters. The task is to detect the appearance of a number embedded in a stream of letters. (b) Example of observed data representing the AB. The graph shows percentage correct answers for the second target (T2) if the first target (T1) has been correctly reported. The attentional blink. If a sequence of distractors is rapidly shown there is a brief period after each correct identification during which the next target is less likely to be seen (Evans et al., 2011, p. 506).

and Somers, 2004). So again, we find that where the eyes are looking and what the focus of attention is can be quite separate.

'You see, but you do not observe. The distinction is clear. For example, you have frequently seen the steps which lead up from the hall to this room.'

'Frequently.'

'How often?'

'Well, some hundreds of times.'

'Then how many are there?'

'How many? I don't know.'

'Quite so! You have not observed. And yet you have seen. That is just my point. Now, I know that there are seventeen steps, because I have both seen and observed.'

(Arthur Conan Doyle, *Sherlock and Watson* in 'A Scandal in Bohemia', 1891)

Generally speaking, as Helmholtz long ago demonstrated, it is perfectly possible to look directly at one object or place, and pay attention somewhere else, a skill now called 'covert attention scanning', as opposed to overt scanning, in which you pay attention where you're looking. You can try this now by keeping your eyes on the page and attending to a location off to one side. Different (though interacting) brain systems seem to be involved in overt and covert orientating of attention: the superior colliculus and frontal eye fields seem to be associated with switches of gaze and attention, while neurons in the posterior parietal cortex are implicated in shifts of attention occurring independently of gaze. The level of interaction is debated, though, with some studies (e.g. De Haan, Morgan, and Rorden, 2008) finding that the brain areas involved in covert and overt attention shifts were virtually identical, but with higher activation levels during overt shifts. This supports the premotor theory of attention (discussed in the next section), which proposes that a covert shift in attention is merely an unexecuted overt shift, using the same neural mechanisms.

These examples show that attention can be involuntarily grabbed as well as deliberately directed, and that attention and gaze sometimes operate together and sometimes not. But this does not necessarily tell us anything about consciousness. We might feel we have consciously chosen where to place our attention without consciousness actually playing a causal role – for example, the feeling of acting consciously might be a by-product or a later effect of the brain processes that selectively direct attention. Returning to our example of the person coming into the room, we might feel as though we experience the sight or sound of the disturbance first and then consciously decide to turn round and look. But whether this is possible may be a matter of timing, a question we return to in Chapter 9. For now, we will briefly survey some of the numerous theories that try to offer a framework for understanding what attention is and how it works. Some of them have their origins in intuitive metaphors or common aspects of experience like turning round when someone opens the door; some focus on the

functional roles of attention; others try to connect attention with other faculties like memory or action. We will start with a brief round-up of the main theories, and then return to the central question of this chapter: how are consciousness and attention related?

THEORIES OF ATTENTION

In the nineteenth century, Helmholtz, Hering, and Wundt were among the physiologists and psychologists who experimented with attention. In the 1950s, many ingenious experiments were performed with a method called dichotic listening, in which two different streams of sound are played to each ear. Normally only one stream can be tracked at once, but certain kinds of stimuli can break through from the non-attended ear, and others can have effects on behaviour without being consciously heard. If the message being listened to moves from one ear to the other, people usually follow the meaning and don't even notice they have swapped ear. This raised the question of whether selection operated early on or after much processing had already taken place, leading to the early versus late selection debate which has never really been resolved – although more recently, as we will see in a moment, it has been sidestepped by the concept of perceptual load.

For a long time most theories treated attention as a bottleneck, with preconscious sensory filters needed to decide what should be let through to the deeper stages of processing (Broadbent, 1958). This makes sense because clearly the brain has a limited capacity for detailed processing, and is a massively parallel system which produces serial outputs, such as speech and sequential actions. So somehow many parallel processes have to be brought together, or selected, to ensure that a sensible serial output occurs.

The main problem with such theories was that to cope with the evidence, the proposed filters became more and more complicated, until the pre-attentive processing began to look as complex as the deeper processing to which it was supposed to give access. These models then gave way to those based on subtler ways of allocating processing resources. The spotlight of attention was then seen as less like a narrow beam or single bottleneck and more like the outcome of many mechanisms by which the nervous system organises its resources, giving more to some items (or features or senses) than others. But for some people, the whole topic was becoming so unwieldy that perhaps the very concept of attention was at fault (Allport, 1993; Pashler, 1998). At some point in its history, the science of attention arguably began to study – or create – something that bears little relation to the intuitive idea of attention as a sharpening of focus (Watzl, 2011). Scientists have redefined attention as a perceptual filter, a feature-binding mechanism, a broadcaster to working memory, or a competitive bias process. Is the gap between the intuitions and the science a problem?

Perceptual load theory was proposed by psychologist Nilli Lavie as an attempt to return to the intuitive idea of a bottleneck of attention, and to rethink it in a simpler way. In this theory, perceptual processing has limited capacity, and when a task involves dealing with a large amount of information (high perceptual load), that capacity is fully exhausted by the processing of the attended-to information: this results in early, top-down selection effects from our current goals and

'Instead of attention having a fixed locus, the [perceptual load] theory argues that awareness depends on the availability of limited-capacity attention'

(Lavie, Beck, and Konstantinou, 2014, p. 8)

priorities. When perceptual load is low, however, spare capacity from processing the task-related information 'spills over', so that we perceive task-irrelevant information via late attentional selection strongly influenced by bottom-up stimuli. This way of thinking saves us from trying to locate a 'fixed locus' of attention, whether acting as a 'gateway' to consciousness or serving some other purpose (Lavie, Beck, and Konstantinou, 2014, p. 8). Nonetheless, it does retain the idea of awareness or consciousness as a location or container which things can get into only if they meet certain attentional criteria. It also relies on the idea of a 'perceptual processing stream from unconscious to conscious levels' (p. 8). Both imply Cartesian materialism (see Chapter 5).

Attention and memory are closely related, and some theories of attention treat short-term memory, with its limited capacity, as the relevant resource to be competed for, or the container to be filled. In other words, being attended to is equivalent to getting into short-term memory, and attention is the same thing as 'the processes that allow information to be encoded in working memory' (Prinz, 2012, p. 93). Other theories do not assume this equivalence, and there are numerous other ways in which attention and memory are connected. In fact, theories of attention are relevant to almost every aspect of brain function, including the neural correlates (Chapter 4), the binding problem (Chapter 6), and unconscious processing (Chapter 8). Predictive-processing models of cognition may also throw light on attention. Studying functional connectivity in large-scale brain networks, psychologist Monica Rosenberg and her colleagues at Yale University (2017) argue that attention is a network property of brain computation and that the functional architecture underlying attention can be measured even when people are not engaged in any task, with individual differences found in the dorsal attention system as well as the default mode network (which is discussed later in this chapter). But here we must concentrate on the core relationship between attention and consciousness.

Neuroscientist Giacomo Rizzolatti and his colleagues in Parma, Italy (Rizzolatti, Riggio, and Sheliga, 1994; Rizzolatti and Craighero, 2010) have suggested a 'pre-motor theory' of selective spatial attention in which attending to a particular position in space is like preparing to look or reach towards it. In experiments using single-cell recording in monkeys, they found that subsets of premotor neurons involved in preparing to make visually guided actions directed towards a particular part of space are selectively activated when attention shifts to that area. Subsequent studies have explored the role of neurons in the frontal eye fields (FEF) within the frontal cortex, finding that stimulating them can elicit saccades and shifts in spatial attention, and that when stimulation is not enough to induce an eye movement, perception is still enhanced for the areas the movement would have been towards. Findings like these highlight the common origin of spatial attention and eye movements, suggesting the existence of a flow of information from visual selection to motor planning, a flow which can be adjusted according to the demands of the task at hand. 'There is no need to postulate two control systems in the brain – one for spatial attention and one for action. The system that controls action is the same that controls what we call spatial attention' (Rizzolatti, Riggio, and Sheliga, 1994, p. 256).

Other experiments have suggested a more complex story, however. FEF seems to contain two separate groups of neurons, one for covert shifts of attention

'The system that controls action is the same that controls what we call spatial attention'

(Rizzolatti, Riggio, and Sheliga, 1994, p. 256)

(without eye movements) and one for overt shifts of attention (with eye movements) (Thompson, Biscoe, and Sato, 2005). That is, counter to what the premotor theory predicts, the FEF neurons driving saccades are separate from those driving attentional selection. Instead, spatially selective activity in the FEF may serve as a visual salience map, identifying potential targets for eye movements without being an explicit saccade plan. The same experiments have also found that when attention shifts covertly to a target in a pop-out visual search task, activity in FEF movement neurons is actively suppressed, with no spatial selectivity. This has led researchers to conclude that activity in the visual, not the motor, FEF neurons is what 'corresponds to the mental spotlight of attention' (2005, p. 9479).

Findings like these suggest that not all areas involved in motor preparation are involved in covert attention, and not all regions involved in covert attention have motor functions. Regions may be involved in both, but in the weaker sense of creating a 'priority map which signals the location of behaviourally relevant stimuli' (Smith and Schenk, 2012, p. 1106). Hence perhaps why dissociations between eye-movement preparation and attention allocation have been found for both overt and covert attention (Hunt and Kingstone, 2003). So the proposal that attention and motor control use the same neural circuits, as well as the stronger claim that motor activation is both necessary and sufficient for spatial attention, may be going too far.

The 'biased competition' (or 'integrated competition') theory, which originated in the 1990s, is one possible alternative. The basic idea is that attention is a neural competition mechanism biased by feedback from a person's goals, expectations, emotional states, and so on (Ruff, 2011). What does this theory say about action control? Here, action preparation increases the probability of the goal of the action being selected for attention and processing, but does not guarantee it, any more than the absence of motor preparation prevents a location being attended to. In this theory, 'attention is the consequence of competition within and across different sensory-motor systems' (Smith and Schenk, 2012, p. 1112). The inputs compete for neural representation, which is allocated on the basis of physical salience, current goals, and working-memory contents, and the winner of the competition is attended to, 'in the sense that it becomes available to higher cognitive processes such as awareness and response systems' (p. 1112). The idea of winning a competition to be broadcast is reminiscent of global workspace's 'fame in the brain' (Chapter 5).

One more theory in the broad cognitive-neuroscience category, and adding an evolutionary spin, is Michael Graziano's 'attention schema theory' of consciousness, which builds on the biased-competition model and on IIT and global workspace theories, and connects consciousness directly with attention. Indeed, for Graziano, 'awareness is the internal model of attention' (Webb and Graziano, 2015, p. 1), and it evolved as a way of modelling and controlling attention: top-down control is improved when the brain can use a simplified model of attention itself. This is the attention schema, and the theory explains 'how the human machine claims to have consciousness and assigns a high degree of certainty to that conclusion' (Graziano, 2016, p. 98). This theory is broadly illusionist, but Graziano prefers to say that consciousness is 'a useful caricature of something real and mechanistic' (p. 112): 'Subjective awareness – consciousness – is the caricature of attention depicted by that internal model' (p. 98).

'attention is the consequence of competition within and across different sensory-motor systems'

(Smith and Schenk, 2012, p. 1112)

'awareness is the internal model of attention'

(Webb and Graziano, 2015, p. 1)

'consciousness is not an illusion but a useful caricature of something real and mechanistic'

(Graziano, 2016, p. 112)



PROFILE 7.1

Michael Graziano (b. 1967)



Michael Graziano is a composer, novelist, and author of children's books as well as Professor of Psychology and Neuroscience at Princeton University. His wide-ranging research includes studies of spatial perception and sensorimotor integration in mon-

keys, how the brain represents the body and its surroundings, and more recently the brain basis of consciousness, including relationships between awareness, attention, and social perception in the human brain. His theoretical work explores the idea that awareness is a construct of the brain's social machinery and his 'attention schema theory' extends this to suggest that awareness is an attention schema computed by an expert system in the brain that attributes awareness to others as well as to oneself. Apart from his surrealist novels, he is author of *Consciousness and the Social Brain* (2010), arguing that awareness is information and consciousness is not mysterious. Graziano is also a skilled ventriloquist when accompanied by his dummy monkey, Kevin.

'attention is the mental activity of structuring the stream of consciousness'

(Watzl, 2011, p. 849)

'Attention is [...] a relevant attribute of the stimulus. It's red, it's round, it's at this location, and it's being attended by me'

(Webb and Graziano, 2015, p. 9)

There are alternatives to accounts which try to reduce attention to neural or computational processes. Instead of focusing on attention's function of selecting items for acting on, we can treat it as something which shapes our experience of the world. In the structuring view of attention, 'attention is contrastive: it structures our mental life so that some things are in the foreground of others', whether or not for the purpose of action selection (Watzl, 2011, p. 849). The conviction that the 'phenomenal character' of attention needs taking seriously if we are to pin down its functional role leads, in this case, back to James's concept of consciousness as a stream: 'attention is the mental activity of structuring the stream of consciousness' (p. 849). In other theories, attention is treated not as a process in itself, but as a manner in which things happen – as an adverb, not a noun. In the cognitive-unison view, for example, the aim is to account for what it takes for someone to perform a task attentively rather than inattentively, and the proposed difference is that the task is carried out with 'cognitive unison' (Mole, 2011). This shifts the question from a *what?* to a *how?*

'[A]ttention is rational-access consciousness', claims philosopher Declan Smithies, in a theory that tries to unify the functional and the phenomenal aspects of attention (2011, p. 268). The idea is that attention is a form of consciousness that makes information fully accessible for use in the rational control of thought and action. 'Rationality' here is a person-level concept: only when high-level processing like reasoning or goal-directed action is based on information that has been attended to is it 'rational' (Smithies, 2011). Here we see that theories of attention are also attempting to characterise how our experience changes when we pay attention, or even equating attention with a kind of consciousness. This takes us back to our initial question of how attention and consciousness relate.

CONSCIOUSNESS AND ATTENTION

There are six main possibilities for how consciousness and attention relate to one another. First, consciousness may depend on attention: we cannot be conscious of something if we aren't paying attention to it. Second, attention may depend on consciousness: we cannot pay attention to something unless we are conscious of it. Third, consciousness and attention may be correlated but not causally connected – maybe because they are both the results of some other mechanism. Fourth, they may be entirely unrelated – in which case the question is why they seem to be related. Fifth, they may actually be the same thing. Or sixth, one or both may be illusory (not be what they seem), or not exist at all – in which case we again have to ask ourselves why we are mistaken.

CAUSAL CONNECTION I: CONSCIOUSNESS DEPENDS ON ATTENTION

The first possibility is that attention is necessary for consciousness: there can be no consciousness without attention.

It is common to feel that we are conscious only or primarily of the things we pay attention to (like being engrossed in a novel I am reading). When they mentioned consciousness at all, the early theories of attention tended to agree, saying that the filters and bottlenecks allowed information ‘into consciousness’, treating it as ‘the sentry at the gate of consciousness’ (Zeman, 2001, p. 1274).

‘information that is not attended cannot reach consciousness’

(Cohen et al., 2012, p. 416)

Today, almost all the processing-based theories of attention propose that conscious experience is the result of attentional mechanisms: the spotlight, filter, perceptual-load, premotor, and working-memory accounts all take this view, as well as the proposal that attention serves to ‘structure our mental life’. Some researchers claim that ‘What is at the focus of our attention enters our consciousness’ (Velmans, 2000, p. 255). Others suggest that ‘*attention* seems to play an especially critical role in determining the contents of consciousness’ (Gray, 2004, p. 166, original italics), or argue that ‘information that is not attended cannot reach consciousness’ (Cohen et al., 2012, p. 416). Dehaene’s neuronal global workspace theory is committed to the view that although considerable processing is possible without attention, attention is required for information to enter consciousness: ‘top-down attentional amplification is the mechanism by which modular processes can be temporarily mobilized and made available to the global workspace, and therefore to consciousness’ (Dehaene and Naccache, 2001, p. 14).

The evidence we looked at in Chapter 3 on inattentive blindness also seems to support this view: if we don’t attend to the gorilla sauntering across the basketball court, we don’t see it. Psychologists Arien Mack and Irvin Rock therefore claim that consciousness depends on attention: ‘there is no *conscious* perception without attention’ (1998, p. 14, original italics). There are other ways of interpreting the findings, however: for example, that what looks like inattentive blindness is actually inattentive agnosia, i.e. we forget having seen the gorilla before we can report it. The most we can say with confidence is that attention seems to be necessary for the kind of consciousness that allows participants to report, after the fact, on the gorilla’s presence.

‘selective attention is necessary, but not sufficient, for a conscious percept to form’

(Koch, 2004, p. 167)

That consciousness is causally dependent on attention also does not mean that attention is solely responsible for shaping consciousness. Here the distinction between necessary and sufficient conditions comes in: attention may be *necessary* to allow or create conscious experience, but not be on its own *sufficient* to do so. This is Benjamin Libet’s view: ‘attention itself is apparently not a sufficient mechanism for awareness’ (2004, p. 115). Christof Koch agrees: ‘selective attention is necessary, but not sufficient, for a conscious percept to form’ (2004, p. 167). Some experiments suggest that we can pay attention (as measured by improved reaction times or response accuracy) to things without being able to report seeing them (e.g. Norman, Heywood, and Kentridge, 2013). We will tackle in depth the problem of distinguishing experimentally between conscious and unconscious responses in Chapter 8.

CAUSAL CONNECTION II: ATTENTION DEPENDS ON CONSCIOUSNESS

Sometimes, however, all this seems backwards. We may often feel that we can consciously direct our own spotlight to pay attention to what we choose. In this sense, maybe consciousness precedes and can direct attention. As James put it: ‘*My experience is what I agree to attend to* [. . .] without selective interest, experience is an utter chaos’ (1890, i, p. 402).

'My experience is what I agree to attend to'

(James, 1890, i, p. 402)

This fits with the feeling that we can consciously choose where to look, which sounds to listen to, or what to think about, and that paying attention can be hard work. James imagines 'one whom we might suppose at a dinner-party resolutely to listen to a neighbor giving him insipid and unwelcome advice in a low voice, whilst all around the guests were loudly laughing and talking about exciting and interesting things' (1890, i, p. 420).

She read, with an eagerness which hardly left her power of comprehension, and from impatience of knowing what the next sentence might bring, was incapable of attending to the sense of the one before her eyes.

(Jane Austen, *Pride and Prejudice*, 1813)

James ultimately came down on this side. His reasons were not scientific; indeed, he concluded that no amount of evidence could really help decide whether consciousness depends on attention or vice versa, and therefore he made his decision on ethical grounds – the decision being to count himself among those who believe in a spiritual force. James was convinced that the essence of volition is 'attention with effort', and that this is central to what we mean by self. So, for him, the answer to this question was vital for thinking about the nature of self and of free will. '*Effort of attention is thus the essential phenomenon of will*', he concluded (1890, ii, p. 562), and by will he meant the genuinely causal force of conscious, personal will. In James's account, consciousness is thought of as a force of will which directs attention; attention then shapes the nature and contents of conscious experience. In a sense, then, this account could fit into the previous category, except that he includes a prior causal stage where conscious willing comes first. Any other theory that puts consciousness first, including most spiritual theories (like those mentioned in Chapter 5), would say the same.

NO CAUSAL LINK BETWEEN CONSCIOUSNESS AND ATTENTION

Moving beyond causal links, there are three possibilities that imply no causal link. First, consciousness and attention might be correlated without being causally connected to each other. Second, they might be altogether distinct. Third, they might in fact be the same thing.

Sometimes awareness of things we are not attending to is an intrinsic and valuable part of our experience, like appreciating the bassline while focusing on the melody: in this case, attention is not even necessary for consciousness. Koch (Koch and Tsuchiya, 2007; Tononi and Koch, 2008) argues that the correlations between consciousness and attention, particularly selective attention, are so patchy and so complex that we must treat them as distinct brain processes; so, consciousness does not reduce to attention. This argument is supported by experiments using a binocular suppression task, which found that activity in V1 is influenced much more strongly by directing attention to a target than by being aware of it (Watanabe et al., 2011). In some cases, awareness and attention even seem to have opposite effects. When the retina adapts to overstimulation, for example, the visual system generates an afterimage, and perceptual suppression (i.e. absence of awareness) makes the afterimage weaker, but so does sustained attention

(Koch and Tsuchiya, 2007). The phenomena of top-down attention without consciousness and consciousness with little or no top-down attention are also not, Koch argues, ‘arcane laboratory curiosities that have little relevance to the real world’ (2007, p. 19): whenever we practise skilled activities that do not require conscious attention, and indeed happen too fast for it, we live this separation. We return to this topic in Chapter 8.

‘top-down attention and consciousness are distinct phenomena that need not occur together’

The other main option in this category is to say the opposite: that consciousness and attention are in fact the same thing. In an integrative account of attention, attention is an emergent property of brain-wide processing – processing which includes the kind of competitive selection posited by the biased-competition account, and may depend on dynamic binding by synchrony (Chapter 6). On this view, the many functions we think of as attentional are manifestations of general processing characteristics in the brain, and ‘there cannot be an anatomically (or functionally) identifiable attentional control system’ (Allport, 2011, p. 27). This theory implies that as soon as we stop trying to claim that attention is causally responsible for consciousness, we may as well say they are the same thing, and so we can refer to attention and to consciousness ‘practically interchangeably’ (p. 49). As such, the phenomena that we call spotlights and bottlenecks and so on are not causal mechanisms, but consequences of those globally integrated neural interactions ‘whose outcome is conscious attention’ (p. 49).

(Koch and Tsuchiya, 2007, p. 16)

‘the many psychological functions generally thought of as attentional [...] reflect general characteristics of the processing network as a whole’

(Allport, 2011, p. 32)

A related option is to say that attention and consciousness are in constant feedback interaction with each other. In Graziano’s attention schema, consciousness is part of the ‘control machinery’ for attention. Awareness tracks attention as its internal model, but when errors creep into the model, attention becomes dissociated from consciousness, and can still operate, but less well. Webb and Graziano (2015) say that the opposite, awareness without attention, is possible – if the internal model wrongly indicated that a perceived stimulus was being attended to – but less likely. In this feedback model, then, attention is the dominant mechanism, and attention and consciousness are separable but normally covary.

ATTENTION AND CONSCIOUSNESS DO NOT EXIST, OR ARE NOT WHAT THEY SEEM

The final possibilities left to consider are that we are so profoundly mistaken about the relation between consciousness and attention that in fact one or both does not exist, or is illusory. Allport’s view leads nicely into these options, by saying that there is no way to separate attention from the rest of the brain in terms of either anatomy or function.

In Chapter 3 we explored the idea that we may be under a grand illusion about consciousness itself, and this will be a thread throughout the rest of the book. When it comes to attention, some researchers are sceptical that attention is a meaningful category at all. There are many reasons to think this: rather than uncovering a coherent set of cognitive or neural mechanisms that can straightforwardly be identified with attention, it becomes ever clearer that attentional processes are diverse and not localised, and most mechanisms involved in attention sometimes operate in the absence of attention. Attention starts to look more like thinking than like perception. Just because we observe lots of attentional effects doesn’t mean there exists *anything* called attention that causes these

effects. Using attention to denote an unspecified causal agent relies on a mysterious homunculus; it is like playing a 'theoretical wildcard' to dodge the need to develop a workable theoretical account (Anderson, 2011, p. 4). If we give up on a unified view, the 'Disunity view' can argue that

Just like chemical analysis shows that jade is not a single kind of mineral (instead there are nephrite and jadeite that are superficially similar),
[. . .] [t]alk about attention does not carve the mind at its joints, because attention, like jade, is not a natural kind.

(Watzl, 2011, p. 848; 2017, p. 32)

'Talk about attention does not carve the mind at its joints, because attention, like jade, is not a natural kind'

(Watzl, 2011, p. 848)

So if we go looking for a specific centre or circuit in the brain responsible for attention, we won't find anything.

Nonetheless, we could still argue that attention is a natural kind at a personal level, even if not at a sub-personal (e.g. a neural) level. The basic argument here would be: if it feels like a meaningful category for talking about conscious experience, that means it is. Applying the same logic would, however, lead us to unquestioningly accept folk concepts like the stream of consciousness or even the soul.

This means that when Sebastian Watzl ties the two together, saying 'Attention is the mental activity of structuring the stream of consciousness' (2011, p. 849), there is a danger of giving reality to things that do not exist. But in any case, as he concludes, studying attention forces us to tackle difficult categories like the differences between states, processes, activities, and manners of going on, and to think carefully about the links between mind and action, and between functional roles and phenomenal qualities, all of which are crucial to thinking about consciousness.

When we start to challenge our intuitions about attention – that there must be a localisable set of brain areas responsible for it, that it is even a unified thing at all – we realise that there is a crucial, profound challenge to be made when it comes to the relation between attention and consciousness. Do we have any way of working out what it means to be conscious of what is being attended to or not being attended to?

The basic problem is that whether and how something forms part of someone's conscious experience can be determined only by either report (what people say) or other explicit decisions (what people do). But reporting on what we see requires us to attend to it. So too do many of the decision-making tasks that are used as criteria for consciousness. So, as philosopher James Stazicker puts it, 'the failure to report an object of visual consciousness might reflect a failure to attend to the object, rather than an absence of visual consciousness of the object' (2011, p. 163). In this case, how can we ever even begin to work out how consciousness and attention relate to or differ from each other? In Stazicker's terms, how could we ever test whether their relationship is one of *dependence* or *independence*: whether the spotlight of attention falling on things is what makes them conscious, or whether it illuminates episodes of consciousness without constituting them?

This problem is another version of the question that Ned Block raises in contrasting phenomenal (P) consciousness and access (A) consciousness. Is there more in conscious experience than can be accessed? There is a long tradition of relevant experiments, beginning with American psychologist George Sperling's

'the failure to report an object of visual consciousness might reflect a failure to attend to the object'

(Stazicker, 2011, p. 163)

experiments in 1960. He showed participants arrays of letters briefly and then cued them to report just one line or column from the array. They could do this accurately, even though they did not know in advance which line or column would be tested, and they could not report all of it. He argued that memory limits the amount that can be reported, and that ‘observers commonly assert that they can see more than they can report’ (Sperling, 1960, p. 26). We might alternatively say that they are conscious of more than they can access, but is this just confusing what should be a simple issue?

Half a century later, Ilja Sligte and colleagues showed that as soon as a stimulus has disappeared, participants can access information from an after-image, and after that can access a limited amount of information from a high-capacity but fragile Very Short Term Memory (VSTM). The concept of VSTM goes back to at least the late 1970s, but Sligte and colleagues locate it to cortical area V4, and conclude that ‘The additional weak VSTM representations remain available for conscious access and report when attention is redirected to them yet are overwritten as soon as new visual stimuli hit the eyes’ (Sligte, Scholte, and Lamme, 2009). Does this mean that the information in V4 is briefly P-conscious and then disappears before it can become A-conscious? Does this confirm a meaningful distinction between the two? And how do we find out? Is a more systematic first-person practice necessary to decide whether the briefly stored information really is phenomenally conscious, or are these third-person studies all that is required to understand what is going on?

Experiments like these have been widely discussed ever since. Block claims that Sperling’s participants had P-consciousness of the specific shapes of all or almost all the letters, but without A-consciousness, that their ‘perceptual consciousness overflows cognitive access’ (Block, 2011).

In any case, we must remember that the relationship between successfully processing or reporting visual information and being conscious of it is far from clear. Here we might turn to what Sperling’s participants said about their experiences – but sadly we have only informal records about this. Block claims that they reported seeing all or almost all the letters. But maybe they were wrong, and thought they saw more than they did. This is the opinion of Stanislas Dehaene and colleagues (2006). It brings us back to the sensorimotor theory of vision we encountered in Chapter 3, and to Dennett’s multiple drafts theory (Chapter 5). In order to answer the question of whether you are conscious of something, you attend to it, which makes you conscious of it, giving you the illusion that you were conscious of it all along, like the fridge light that is always on when you look.

But maybe we are making unnecessarily complicated assumptions about people’s illusions about their own experience. Maybe we can take people’s reports about their consciousness at face value, Stazicker suggests. Maybe they were conscious of all the letters, but not of the specific shapes of all the letters. Maybe, then, their reports exactly match their experiences. On this account, what the cueing did was not make accessible some portion of an already conscious experience; instead, it made more determinate some information in that experience. This is not the same as claiming that, as in multiple drafts, attention exerts retroactive effects on our conscious experience, or what we think of as that experience. In this intermediate view, attention exerts effects on consciousness as it happens.




FIGURE 7.4 • Sperling (1960) showed arrays like this very briefly and then cued participants to report a single line or column.

• SECTION THREE: BODY AND WORLD

What leads us astray is perhaps the tendency to assume that vision is always maximally clear (Stazicker, 2011) – and we saw in Chapter 3 that this is not the case, even in a basic sense like how much the resolution of the retina diminishes towards the periphery. It may also be quite plausible that the moment at which we gain access to more detail would go unnoticed – it might well not be something people would, or could, report. Paying attention might so naturally give us access to more detail that we would not even remark on the change. The findings from Carrasco and colleagues which we mentioned earlier add weight to the idea that attention changes the specifically visual quality (something as specific as the spatial resolution) of visual consciousness: accessing something by paying attention changes the *quality* of the experience rather than just *whether* we experience something.

We can ask similar questions about Mack and Rock's experiments used to demonstrate inattention blindness. Can we necessarily assume that people did not report seeing the additional stimulus because they were unconscious of it? Along similar lines to the discussion of Sperling's experiments, failure to report seeing an additional stimulus or to identify it

might reflect either (i) that the subjects were not visually conscious of the stimulus, or (ii) that, though subjects were conscious of the stimulus, they did not attend to it in the way required for this consciousness to form the basis for a reliable decision. To assume that (i) is the correct interpretation is to beg the question. On the other hand, there's no obvious way to argue for interpretation (ii) either, because without reports or fairly explicit decisions we lack compelling evidence for the presence of consciousness.

(Stazicker, 2011, p. 164)

It is easy to assume that consciousness is all-or-nothing, on or off: we are conscious of something or we are not. But this may be one of the errors that



FIGURE 7.5 • Meditating in Japan's famous rock garden at Ryoanji.

prevents us from accurately assessing its relationship to attention. So where does this leave us? Searching for the neural correlates of consciousness might still sound like a nice idea: if we could find out what neural activity correlates with (say) visual consciousness, we could determine whether this activity ever occurs without those processes which correlate with attention. But this takes us back to problems we encountered in Chapter 7: how do we establish these correlates without first knowing whether or not one occurs without the other? Some people, like the philosopher Hilary Putnam, conclude that there is simply no way of answering the question of whether there is unreportable consciousness.

MEDITATION AND ATTENTION

The science of consciousness often feels like a science of unanswerable questions. Every scientific experiment we discuss in this book tries, in one way or another, to negotiate the boundary between measurable, objective facts that may or may not be relevant to consciousness, and the subjective reality of consciousness: the *what it's like to be*. Later in the book we will devote more time to looking at alternatives to the 'third-person' science that is standard in most of today's experimental psychology and neuroscience labs: ways to bring the first and the second person (me and you) into scientific practice. For now, in the last part of this chapter, we will focus on one aspect of this more inclusive science – one inspired by the ultimate first-person practice: meditation.

Meditation is also the ultimate training of attention. There many different forms of meditation, but the first step in nearly all of them is calming the mind. This skill can take many years to master, but then it becomes easy to sit down and let the mind settle. Everything that arises is let go, like writing on water. Nothing is met with judgement or opinions, and as reactions gradually cease, clarity appears. The sounds of birds, the sight of the floor, the itch on the hand, they are just as they are: suchness. Many traditions claim that in this decluttered state, insight into the mind can spontaneously arise.

Those who practise certain kinds of meditation claim that they awake from illusion and see directly the nature of mind. If they are right, their claims are important both for the introspective methods they use, and for what they say about consciousness. But are they right?

Many interesting questions are posed by these practices. For now, we will focus on their relevance to attention. To start with, we will briefly sketch out what meditation is, how it is done, and what its effects are.

MOTIVATIONS AND METHODS

Most methods of meditation have religious origins. In particular, Buddhism, Hinduism, and Sufism have long traditions of disciplined meditation, but comparable methods of silent contemplation are found within the mystical traditions of Christianity, Judaism, and Islam (Ornstein, 1986; West, 1987). Within these traditions, people meditate for widely different reasons. The more religious may want to gain merit, get to heaven, or ensure a favourable reincarnation, while others sit for insight, awakening, or enlightenment.

Many secular methods have emerged from religious traditions. For example, in *Buddhism without Beliefs* (1997) and *After Buddhism* (2015), Stephen Batchelor maps out ways of practising with no religious connotations or commitment to belief in gods, persisting selves, or life beyond death. Transcendental Meditation (TM) was derived from Hindu techniques, brought to the West in the late 1960s by Maharishi Mahesh Yogi, and is now taught within a large, hierarchical, and highly profitable organisation which claims that TM provides deep relaxation and inner happiness, eliminates stress, and improves relationships, sleep, health, creativity, efficiency, concentration, confidence, and energy. (See our website for material on the wilder claims made for it.)

[Mindfulness is] the awareness that emerges through paying attention on purpose, in the present moment, and nonjudgmentally to the unfolding of experience moment by moment'

(Kabat-Zinn, 2003, p. 145)

Mindfulness lies at the heart of most meditation methods and can be practised in the rest of life as well as when sitting in meditation. It is usually defined in terms of an acceptant, nonjudgemental focus on the present moment, without discrimination, categorisation, judgement, or commentary. It is 'the active maximising of the breadth and clarity of awareness' (Mikulas, 2007, p. 15). Mindfulness-based stress reduction (MBSR) is a technique developed by Jon Kabat-Zinn at the University of Massachusetts Medical Center in the 1970s and now used widely for conditions ranging from depression and anxiety to pain management and heart disease. He defines mindfulness as 'the awareness that emerges through paying attention on purpose, in the present moment, and nonjudgmentally to the unfolding of experience moment by moment' (2003, p. 145). This is surprisingly hard to achieve more than very briefly. Training is usually for eight weeks and includes a mixture of mindfulness meditation and yoga.

Despite their different origins, the basics of all types of meditation might be summed up in the words 'pay attention and don't think'. It is hard to believe that such a simple practice could create the kinds of transformations and insights claimed by some meditators, yet this is essentially the task undertaken. It is surprisingly difficult, as you will know if you have tried, and the many varieties of meditation can be seen as different ways of easing the task. If you have never tried it, just take ten seconds now and see whether you can *not think* for that short length of time.

Meditation usually involves sitting in a special posture, such as the full lotus or other less strenuous crossed-legged positions (see the website for more detail), but there is nothing mysterious about this. The postures all serve to keep the body both alert and relaxed, while keeping still for long periods. It is possible to meditate in any position at all, and TM suggests just sitting comfortably in a chair, but the two main dangers are becoming too tense and agitated, or falling asleep. The traditional postures help to avoid both, as well as encouraging good breathing and an upright spine.

During long meditation retreats, sitting is sometimes alternated with very slow walking meditation, or even fast walking or running meditations, to provide some exercise and stimulation without disturbing the practice. In fact, for some traditions the ultimate aim is to integrate meditation into all life's activities.




FIGURE 7.6 • Traditional meditation postures all achieve a stable and comfortable position with an upright spine, to encourage a state of alert relaxation. Sitting on a low bench achieves the same objective and is more comfortable for those not used to sitting on the floor.

She then became haunted by a suspicion which she was so reluctant to face that she welcomed a trip and stumble over the grass because thus her attention was dispersed, but in a second it had collected itself

again. Unconsciously she had been walking faster and faster, her body trying to outrun her mind; but she was now on the summit of a little hillock of earth which rose above the river and displayed the valley. She was no longer able to juggle with several ideas, but must deal with the most persistent, and a kind of melancholy replaced her

excitement. She sank down on to the earth clasping her knees together, and looking blankly in front of her. For some time she observed a great yellow butterfly, which was opening and closing its wings very slowly on a little flat stone.

'What is it to be in love?' she demanded, after a long silence; each word as it came into being seemed to shove itself out into an unknown sea. Hypnotised by the wings of the butterfly, and awed by the discovery of a terrible possibility in life, she sat for some time longer. When the butterfly flew away, she rose, and with her two books beneath her arm returned home again, much as a soldier prepared for battle.

(Virginia Woolf, *The Voyage Out*, 1915)

BASIC PRINCIPLES

Common to all forms of meditation are two basic tasks: paying attention and not thinking. Both raise interesting practical and theoretical questions. What do you pay attention to? How do you maintain concentration? How do you not think? The different methods outlined below give different answers, but almost all techniques share common methods for dealing with unwanted thoughts.

Pushing thoughts away does not work, as Daniel Wegner showed when he asked people not to think about a white bear. Not only did they fail, but thought suppression could even lead to obsession (Wegner, 1989). Unwanted thoughts may be held at bay temporarily, but then they come back with greater force, or change into other more persistent thoughts, or set up emotional states that keep reigniting them. The answer is not to fight against thoughts but to learn to let them go and return to the practice. If you get angry with yourself for being so easily distracted, just let the anger go, too.

Despite many differences, all meditation methods can be conveniently divided into two main types: open or nondirective versus concentrative (Ornstein, 1986; Wallace and Fisher, 1991; Farthing, 1992; Xu et al., 2014), receptive versus concentrative (Austin, 2009), or open monitoring versus focused attention (Lutz et al.,



ACTIVITY 7.1

Meditation

Meditation can be done by yourself or in a group. First, sit down comfortably. You should have your back upright but be able to relax, with your head floating lightly at the top of your spine. If you know how to sit in a meditation position, do so. If you wish to try one, make sure the floor is not too hard or use a rug or blanket, and choose a firm cushion to sit on. Cross your legs in the way that is easiest for you and make sure that you can keep your back upright without pain. Otherwise, sit upright in a straight chair with your feet flat on the floor and your hands gently resting on your lap. Look at the floor about two feet in front of you, but don't concentrate hard on one spot, just let your gaze rest there gently. If it wanders, bring it back to the same place.

Set a timer to ten minutes.

Begin by just watching your breath as it flows in and out. When you are ready, begin counting. On the first out-breath, count 'one' silently, and then on the next out-breath 'two', and so on. When you get to ten, start again at one, and continue until the timer sounds. That's all.

Your attitude towards everything that arises should be the same: 'Let it come, let it be, let it go.' When you realise that you have slipped into a train of thought, just let it go and return to watching your breath and counting. Do not fight the thoughts or try to force them to stop. Just let go. Do the same with sounds or sights or bodily sensations: just let them be. This way they won't be distracting at all.

Just one session may show you something about your own mind. If you wish to do more, commit yourself to meditating every single day for a week, perhaps first thing in the morning, or twice a day if you think you can manage it. It is better to sit for ten minutes every day without fail than to try to do more and give up.

'the long path of Zen involves a "letting go"'

(Austin, 2009, p. 48)

2008; Lippelt, Hommel, and Colzato, 2014). Some schemes distinguish active from passive techniques (Newberg and D'Aquili, 2001), but, as we shall see, there are active and passive aspects to them all. Sometimes both methods are used within the same tradition or even within the same session for different purposes. Paying attention single-mindedly for long periods is not easy, and the various methods can be seen as different ways of making it possible.

OPEN MEDITATION

Open or receptive meditation means paying attention equally to everything that is happening, whether that is perceptions, feelings, or thoughts, but without responding. This is usually done with the eyes open or half-open.

Mindfulness meditation is a form of open meditation derived from Buddhism, and in particular from the method of *shikantaza*, which means 'just sitting'. With practice, this deliberate, present, nonjudgemental openness to everything leads to what is called choiceless awareness, bare awareness, or bare attention.

One of the first effects that new meditators notice is how different this is from their normal state of mind. They have 'the piercing realization of just how disconnected humans normally are from their very experience' (Varela, Thompson, and Rosch, 1991, p. 25). You will not, however, be so 'disconnected' if you have already been doing the practices in this book. Although not obvious at the start, you may now see how they have been building up your attentional skills and making you more familiar with your own mind. Indeed, even the first question, 'Am I conscious now?', is a way of bringing yourself out of distraction and into mindfulness. We hope this process will continue as you read the rest of the book.

Mindfulness is a direct and simple technique, but difficult to do. When thoughts and distractions arise, the task is to return to the present moment, but this is not easy when the present moment is full of pain in the legs, memories of unhappiness, anger at yourself or someone else, or anticipations of future pleasure. One solution is to meet all these distractions with the attitude 'Let it come, let it be, let it go'. These are the three reminders given to Zen students by British psychologist and Zen master John Crook (1990).

'let thru – let be – let go'

(Crook, 1990, p. 160)

'Let it come' means letting a thought, feeling, or perception arise without trying to block it. 'Let it be' means not reacting to it, trying to get rid of it, or judging it good or bad. 'Let it go' means letting it come to its natural end without either holding it back or engaging with it. Although mindfulness is primarily a meditation technique, it can be practised at all times, and for some Buddhists the aim is to remain fully present in every action and every moment of waking life, and even during sleep. This means never giving in to distraction or desire, never dwelling in the past or future, and being attentively open to everything, all the time. This is a radically different way of living.

CONCENTRATIVE MEDITATION

Concentrative meditation means paying focused attention to one thing without distraction, rather than remaining open to the wider world. In a famous study in the early 1960s, American psychiatrist Arthur Deikman rounded up a group of friends, sat them in front of a blue vase, and asked them to concentrate on it for half an hour, excluding all other thoughts, perceptions, and distractions. The

effects were very striking. The vase seemed more vivid, richer, or even luminous. It became animated, or alive, and people felt they were merging with the vase or that perceived changes in its shape were happening in their own bodies. This sounds like an intensified version of the increases in contrast or resolution that may result from everyday attending. Deikman argued that as we normally develop through life, we learn to attend increasingly to thoughts and abstract categorisations. This allows us to conserve attentional energy for the higher-level goals of biological and psychological survival. But the side-effect is that our perception becomes automatised and dull. The effect of this exercise in concentrated meditative attention was 'deautomatisation' (Deikman, 1966, 2000). Similar effects may be observed after taking LSD or in other 'altered' states of consciousness (see Chapter 13). The mechanisms involved in reducing automatised and increasing cognitive-emotional flexibility through meditation may include: reducing the chaining of thoughts into an associative stream; making the contents of the thought chains more flexible and varied; and/or creating new paths for the chains of thought (Fox et al., 2016).



FIGURE 7.7 • Letting go, not pushing away.

The most common object for concentrated attention is the breath. One method is to count out-breaths up to ten and then start again at one. This can help deal with the distractions that all too often plague the mind in open meditation, leading one to get lost in long trains of thoughts for minutes at a time. If you are counting the breath, you are much more likely to notice that you have lost count, and even to remember where you got to. This can be quite shocking as well as useful. Another method is just to watch and feel the sensation of air flowing naturally in and out as the chest rises and falls.

Sometimes special techniques are used that alter the breathing rate or depth, the ratio of in-breath to out-breath, and whether the breathing is predominantly in the chest or abdomen. Different breathing patterns have powerful effects on awareness and there is evidence that experienced meditators use these effects instinctively. For example, during the in-breath, pupils dilate, heart rate increases, and activity in the brain stem increases, as does activity in some higher brain areas. The opposite occurs during the out-breath. Blood gas levels also change. Research shows that experienced meditators spend more time slowly exhaling, and increase abdominal breathing. Overall, they may reduce their breathing rate from the normal twelve to twenty or so breaths per minute to as little as four to six, often without ever explicitly being trained or even realising that they are doing so (Austin, 1998). Some meditators stop breathing altogether for periods of many seconds, and one study of TM adepts showed that these stops often coincided with moments of 'pure consciousness' or wakeful no-thought (Farrow and Hebert, 1982; Forman, 1990, 1999; see Chapter 18).

• SECTION THREE: BODY AND WORLD

Mantras are words, phrases, or sounds repeated either silently or out loud. When thoughts arise, the meditator just returns attention to the mantra. Mantras are used in Buddhism, Judaism, and Hindu yoga, including the well-known *Om Mani Padme Hum*, which means ‘the jewel in the centre of the lotus’. In Christianity, the early Desert Fathers used to repeat *kyrie eleison* (from the Greek for ‘Lord, have mercy’) silently to help them achieve a state of ‘nowhereness and nomindness’ (West, 1987). *The Cloud of Unknowing* recommends clasping a word such as *God* or *love* to your heart so that it never leaves, and beating with it upon the cloud and the darkness, striking down thoughts of every kind, and driving them beneath the cloud of forgetting, so as to find God and achieve complete self-forgetfulness (14th century/2009, pp. 24–25).

TM is based on mantras, and new students are given a ‘personal’ mantra. In fact, this is assigned merely on the basis of age, and probably the words do not matter. Indeed, *anything* can be used as a focus of attention, and common ‘aids’ include candle flames, flowers, stones, or any small object. Some traditions use mental images, which range from simple visions of light to the highly elaborate sequences of visualisation taught in Tibetan Buddhism.

Finally, in Zen Buddhism, and in particular in the Rinzai school, practitioners concentrate on koans or hua tous. These are questions or short stories designed to challenge the intellectual mind with paradox, polarity, and ambiguity, and force it into a state of open inquiry. Some meditators use the same koan for a whole lifetime, such as the question ‘Who am I?’, or the question ‘What is this?’ used in Korean Zen (Batchelor, 2001; and see Concept 18.1 for other examples). Others pass through a series of koans as they develop their understanding. Koans are designed not to be answered but to be used.

How does this bewildering array of methods, all called ‘meditation’, help us to understand consciousness? The effects of meditation on physical and mental health may be relevant to this question, and we will come back to them in Chapter 13, but the implications can go deeper still. After long practice with any of these methods, meditators claim that letting go gets easier, and thoughts and feelings that would previously have been distracting become just more stuff appearing and disappearing without response. Ultimately, in alert and mindful awareness, the differences between self and other, mind and contents, simply drop away. This is known as realising nonduality.

Can this really be possible? The hard problem confronts us precisely because of these same dualities. So the suggestion that it is possible to transcend them should be of great interest indeed to a science of consciousness. Attention plays a crucial role in this possible transcendence.

‘Can I catch myself not attending, without attending?’

(koan on attention)

ATTENTION, ATTENTION, ATTENTION

A man asked the fourteenth-century Zen master Ikkyu to write for him some maxims of the highest wisdom. Ikkyu wrote ‘Attention.’ Dissatisfied with this answer, the man asked for more. He wrote ‘Attention. Attention.’ The man complained that he saw nothing of much depth or subtlety in that. So Ikkyu wrote ‘Attention. Attention. Attention.’ When the man angrily demanded to know what attention means, Ikkyu gently answered ‘Attention means attention’ (Kapleau, 1980).

You may be irritated by such stories. You may think that the way the brain directs its resources has nothing to do with wisdom. Yet there have been many studies of the effects of Zen and other meditative training on brain function, and it seems that the development of attention may be critical.

The very earliest studies began back in the 1950s when intrepid researchers carried cumbersome EEG equipment up to the monasteries and mountain caves of Indian yogis, and recorded brain waves while banging cymbals, flashing lights, and plunging the yogis' feet in cold water (Bagchi and Wenger, 1957). The yogis were not distracted by these violent intrusions, and for a while this seemed to confirm the difference between their concentrative meditation and the open meditation of a group of Japanese adepts who appeared to remain alert to sounds and lights with no sign of habituation (Kasamatsu and Hirai, 1966). Sadly, this simple picture was not confirmed by the conflicting results and hypotheses that followed (Fenwick, 1987; West, 1987), and it was many years before research into the neuroscience of meditation began to make progress again, with discoveries concerning changes in functional connectivity, shifts in attention, and changes in the default mode network.

James Austin is an American neurologist who undertook extensive Zen training in Japan and has since explored the relationships between Zen and the brain. Reviewing numerous studies, he concluded that the two main types of meditation, concentrative and receptive, differentially train one or other of the two main attentional systems in the brain: the dorsal system that mediates purposeful, voluntary, or high-level attention, and the ventral attention system that controls vigilance, alerting, and involuntary attention (Austin, 1998, 2009). Furthermore, he argues that the top-down skills of the dorsal system are relatively easy to acquire and can be seen in short-term studies with novice meditators, while 'advanced meditators may slowly be developing more "opening-up" meditative styles and engaging in a range of subtle, global, more bottom-up receptive practices' (Austin, 2009, p. 43).

Similar conclusions are drawn by Antoine Lutz from his neurophenomenological research at the University of Wisconsin-Madison. His group found that focused meditation involves training the neural systems associated with monitoring conflict (e.g. the dorsal anterior cingulate cortex and dorsolateral prefrontal cortex), paying selective attention (e.g. the temporal-parietal junction, ventro-lateral prefrontal cortex, frontal eye fields, and intraparietal sulcus), and sustaining attention (e.g. right frontal and parietal areas and the thalamus). By contrast, open meditation does not involve an explicit attentional focus, and so should rely on brain regions implicated in monitoring, vigilance, and disengaging attention from distracting stimuli (Lutz et al., 2008).

Later experiments compared mindful self-awareness (with emotion- and sensation-based prompts, like 'feel into yourself') against self-referential thinking (with cues to thinking: 'reflect who you are') in novice and experienced meditators. For the self-awareness group, they found deactivation of prefrontal and precuneus areas associated with mind-wandering and the default mode network, especially in long-term meditators, and in both groups there was greater activation of areas associated with somatosensory attention (J. Lutz, 2016, pp. 21–34). Many subsequent findings support this distinction, confirming the proposal that focused and

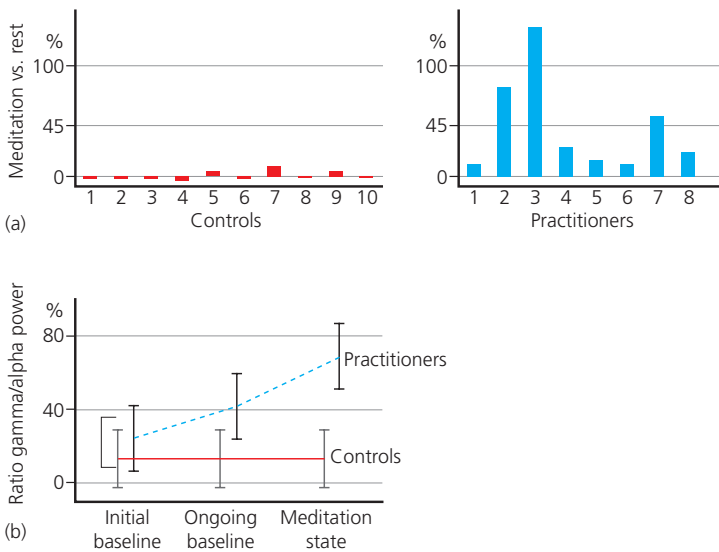


FIGURE 7.8 • (a) The ratio of gamma (25-42 Hz) to slow oscillations (4-13 Hz) averaged within individuals across all electrodes during compassion meditation. The X-axis represents the participant numbers; the Y-axis represents the difference in the mean ratio between the initial state and meditative state. (b) The significant interaction between group (practitioner, control) and state (initial baseline, ongoing baseline, and meditation state) for this ratio. The relative gamma increase during meditation was higher in the post-meditation session. In the initial baseline, the relative gamma was already higher for the practitioners than the controls and correlated with the length of the long-term practitioners' meditation training through life (from Lutz et al., 2008).

open meditation bias processing towards goal persistence and cognitive flexibility respectively (e.g. Colzato et al., 2016).

More generally, there is accumulating evidence that meditation and self-reported mindfulness enhance processing speed and cognitive flexibility and reduce susceptibility to cognitive interference (Moore and Malinowski, 2009). For example, German psychologist Peter Malinowski (2013) reviewed evidence that meditation training leads to attentional resources being used more efficiently and flexibly. Many studies on meditation and attention use the Stroop word-colour task, where people have to name the colour of the font a word is presented in. Normally, when there is a conflict (e.g. the word 'red' is presented in green), people

are slower to respond, or make more errors if forced to respond quickly. Meditators have been found to perform significantly better than non-meditators on this task, and even short mindfulness training sessions can make a difference. This suggests that meditation reduces automaticity and thus improves attentional control. In one study, participants were given a raisin to eat. Those who received brief guidance in how to monitor their sensory experiences of eating it were better at detecting unexpected distractors during a goal-directed task (Schofield, Creswell, and Denson, 2015). Interestingly, although mindfulness is often suggested to help combat the effects of stress, in this experiment it did not offer any protective effect against 'cognitive depletion'. Regardless of the raisin-eating training, being cognitively depleted by a difficult writing task made people more likely to process the perceptual details of the distractor rather than returning to the task at hand.

Meditation research also taps into a growing interest in the activities of the resting mind. In 2001, neurologist Marcus Raichle and his colleagues were conducting experiments that required participants to concentrate on a demanding task. Such experiments commonly use a resting condition as a control, assuming this is uninteresting, but Raichle noticed that during the task, activity in certain areas of the cortex was reduced and then increased again between tasks. There seemed to be 'an organized, baseline default mode of brain function that is suspended during specific goal-directed behaviors' (2001, p. 676). This is how they accidentally discovered what they termed the 'default mode'.

The default mode network (DMN) is active when someone is awake but not focused on a specific task, such as during mind-wandering or daydreaming. Mind-wandering, or 'task-unrelated thought', tends to include thinking about oneself or other people, and remembering the past or imagining the future. Activity in this network is negatively correlated with other brain networks, especially those involved in focused attention. Major DMN hubs include the posterior cingulate cortex (PCC) and the medial prefrontal cortex (mPFC). The PCC is thought to be involved in continuous broad-based sampling of external and internal environments when focused attention for task-specific activity is not required. The mPFC plays a role in mediating the visceral and motor aspects of emotional information,

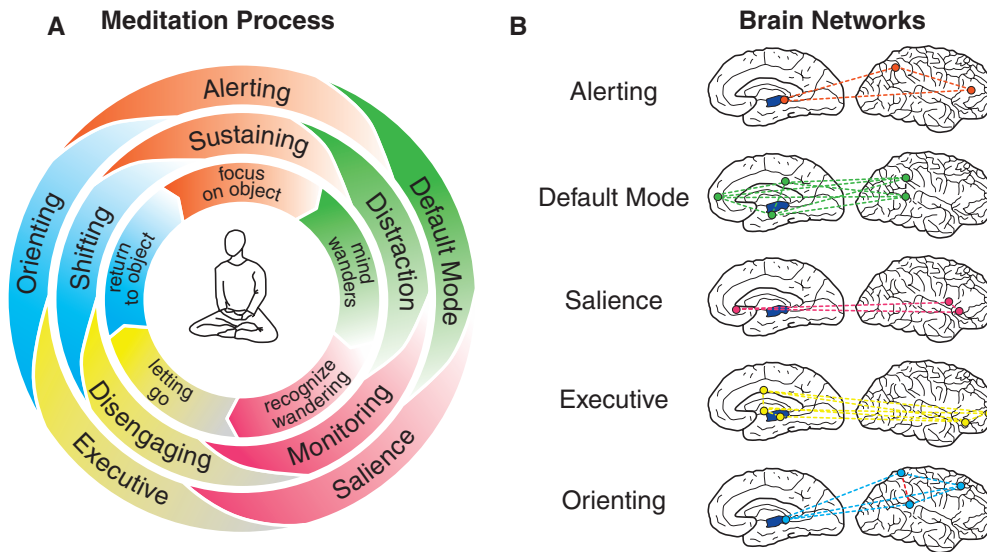


FIGURE 7.9 • Effortful attention regulation during meditation. Panel (a) provides a schematic representation of the meditation process. The inner circle outlines the phenomenological layer, presenting the typical sequence (clockwise) a meditator will go through. The middle circle relates the attentional processes that lie underneath, while the outer circle represents the different brain networks that are involved in carrying out these functions. The different attentional processes and the brain networks are represented as partially overlapping to indicate that in many instances more than one process/network is involved. Panel (b) outlines the main brain areas involved in each of the five networks. Mindfulness training also seems to reduce inattention blindness (Malinowski, 2013, p. 4).

and, like the PCC, is associated with introspective processing which diminishes when attention is outwardly directed (Broyd et al., 2009). The DMN strengthens during normal human development, but various forms of meditation can reduce its activation and its connectivity, including by bringing about structural changes in areas like the PCC, temporoparietal junction, and precuneus (Brewer et al., 2011; Fox et al., 2014). Conversely, nondirective meditation, in which mind-wandering is encouraged, shows enhanced activation of the DMN (Xu et al., 2014).

The concept of the DMN has made ‘mind-wandering’ a fashionable topic, so we now know that the interference between primary task demands and mind-wandering is not as straightforward as it first appeared (Thomson, Besner, and Smilek, 2013), and that mind-wandering seems to be unique in involving cooperation of both default and executive network regions, which are usually opposed (Christoff et al., 2009). Another concept that has grown in importance is interoceptive attention, or attention to bodily sensations related to digestion, blood flow, breathing, and proprioception, which has been proposed as crucial to mindfulness meditation. One study used brain imaging to try to distinguish meditators from non-meditators on the basis of subtly different patterns of activation. This succeeded with thirty-seven out of thirty-nine participants (Sato et al., 2012). The most informative areas in making this distinction were involved in awareness and recognition of bodily sensations (see also Manuello et al., 2016).

Together, findings from this line of research suggest that meditation can bring about rather profound changes in the brain’s global responses to the body and the world, and that what we call attention must be crucial to how it does so.

Could this research go further and begin to build bridges across the explanatory gap or the great chasm? The Churchlands argue that in past centuries people

‘Responsiveness to the world, in action, precisely involves a way of attending to the world, more often unconscious than not’

(Wu, 2011, p. 112)

'subjective intuition [about selective attention] does not coincide with and is, in fact, contradicted by experimental evidence'

(Rizzolatti et al., 1994, p. 231)

'no one knows what attention is, and [...] there may even not be an "it" there to be known about'

(Pashler, 1998, p. 1)

struggled to understand that light really *is* electromagnetic waves or that heat *is* kinetic energy, but now it seems obvious. Today's students already think of depression, addiction, and learning as changes in brain state, they say, and as we older folk die off, people will come to accept that subjective experience just *is* a pattern of brain activity (in Blackmore, 2005). So, could future meditators who shift from distraction into open awareness imagine, or even *feel*, this to be the dropping out of the dorsal attention system and engagement of the ventral allocentric system? If so, subjective and objective would not seem so far apart. But perhaps the objective side of the equation will have to involve more than brain activity, taking in the rest of the body and the world as well.

Attention is central to these brain–mind–world connections, and changes in attention may be amongst the most profound reasons why people put so much time and effort into learning to meditate. This is not in order to gain something measurable, nor to achieve a temporary state of consciousness, nor to reduce stress. Instead, it may be about coming to perceive and be in the world differently. They may be prepared to face new stresses and much hard work in order to see through some of the common illusions about consciousness, and wake up.

In this chapter, we have considered numerous theories of what attention is and does. We have seen that attention can be involuntarily grabbed or deliberately controlled, and that both are intimately related to eye movements and preparation for action; we have looked at the effects of systematically training our powers of attention; and we have asked whether paying attention changes the quality of conscious experience, and whether we can ever really establish where the dividing line is between consciousness and attention. But all this still leaves us with one of our early questions about the role of consciousness.

The fact that we may feel we have consciously chosen where to place our attention does not necessarily mean consciousness actually does play a causal role: for example, the feeling of acting consciously might be a by-product or later effect of the brain processes that selectively direct attention. Returning to our example of the person coming into the room, you might feel as though you experience the sight or sound of the disturbance first and then consciously decide to turn round and look, but we pointed out that the sequence is likely to be the other way around. In Chapter 9, we will meet the idea that consciousness takes some time to build up, and what, if anything, this means for our sense of agency. But first, in Chapter 8 we will explore the difference between doing something consciously and doing it unconsciously.

READING

Lavie, N. (2007). Attention and consciousness. In M. Velmans and S. Schneider (Eds), *The Blackwell companion to consciousness* (pp. 489–503). Oxford: Blackwell.

An overview of the early/late selection debate and perceptual load theory, emphasising direct versus indirect measures, with illustrations from change blindness and inattention blindness.

Malinowski, P. (2013). Neural mechanisms of attentional control in mindfulness meditation. *Frontiers in Neuroscience*, 7, article 8.

The neuroscience of how meditation affects attention.

Manuello, J., Vercelli, U., Nani, A., Costa, T., and Cauda, F. (2016). Mindfulness meditation and consciousness: An integrative neuroscientific perspective. *Consciousness and Cognition*, 40, 67–78.

The present and future of the (neuro)scientific study of consciousness and meditation.

Shear, J., and Jevning, R. (1999). Pure consciousness: Scientific exploration of meditation techniques. *Journal of Consciousness Studies*, 6(2–3), 189–210. Response: Nixon, G. (1999). A ‘hermeneutic objection’: Language and the inner view. *JCS*, 6(2–3), 257–269 (also includes response to response: Shear, J. (1999). Reply to Nixon on meditation). Reprinted in Varela and Shear (1999), pp. 189–209, 257–269.


A debate on what, if anything, meditation can offer to the study of consciousness.

Stazicker, J. (2011). Attention, visual consciousness and indeterminacy. *Mind & Language*, 26(2), 156–184.

Philosophical criticism of confusions about attention and visual consciousness (including Block’s P/A distinction), arguing that visual indeterminacy helps us think about the role of attention.

Watzl, S. (2011). The nature of attention. *Philosophical Compass*, 6(11), 842–853.

Categorises theories of attention, from intuitive conceptions to scientific accounts. (A companion piece, ‘The philosophical significance of attention’ (*Philosophical Compass*, 10, 722–733, expands on the relation between attention and consciousness.)



Conscious and unconscious

CHAPTER

EIGHT

'The power of the unconscious' is a common phrase reflecting the popular notion that our minds are divided in two. Not only do we instinctively separate mind from body, we split mind itself into parts. Amazing powers are sometimes attributed to 'the unconscious', while the conscious mind is derided as more rational and restricted. We may be urged to unleash our unconscious potential or listen to what bubbles up from the depths of our unconscious minds. The opposite happens too: we may feel that what makes us human is our ability to overrule our animal instincts with reason, or not to let our emotions get the better of us.

The distinction between conscious and unconscious is often likened to that between mind and body: we tend to talk about conscious processes, into which we feel we have full insight, as mental ones ('I've given this a lot of thought', 'I know that this is a bad move') and about unconscious processes, which remain opaque to us, as embodied ('I have a gut feeling about this', 'He makes my skin crawl', 'The whole idea just feels wrong').

The idea that the mind is divided into parts can be traced back as far as early Hindu texts or ancient Egyptian beliefs about sleep and dreams, and to Plato, who gave the soul three parts: reason, spirit, and appetite, all with their own goals and abilities (Frankish and Evans, 2009). The idea appears again in eighteenth-century Western philosophy, in Western literature, for example in Shakespeare and Coleridge, and in early twentieth-century psychoanalysis. Typically, the 'highest' faculty (e.g. reason) is thought of as separate from the body, while instinct is understood as a base, bodily function which connects us with other animals.

*O the mind, mind has mountains; cliffs of fall
Frightful, sheer, no-man-fathomed. Hold them cheap
May who ne'er hung there. Nor does long our small
Durance deal with that steep or deep. Here! creep,
Wretch, under a comfort serves in a whirlwind: all
Life death does end and each day dies with sleep.*

(Gerald Manley Hopkins, from 'No worst, there is none', c. 1885)

Gradually the focus shifted from parts of mind to mechanisms, and to distinct types of processing going on in one brain. This can be traced back at least to Helmholtz's idea of 'unconscious inference', to William James's distinction between associative and true reasoning, and more recently to the debates over subliminal perception, unconscious processing, and 'dual-process theories'. These theories can come in many forms, applying to memory, learning, or decision-making, but most suggest that one process is fast, automatic, inflexible, effortless, and dependent on context, while the other is slow, effortful, controlled, flexible, requires working memory, and is independent of context. The two kinds of process map easily onto a distinction between unconscious and conscious processes.

The fact that similar distinctions have been rediscovered or reinvented throughout the history of philosophy and psychology leads some to believe that 'this reflects on the nature of the object of study that all these authors have in common: the human mind' (Frankish and Evans, 2009, p. 2). In other words, the distinction is common because it is valid.

Certainly it is often taken for granted in consciousness studies. For example, an encyclopaedia entry on the 'contents of consciousness' begins: 'Of all the mental states that humans have, only some of them are conscious states. Of all the information processed by humans [. . .], only some of it is processed consciously' (Siegel, 2009, p. 189). It seems so obvious!

But is this right? Or could it be another example of powerful intuitions leading us astray? An alternative is that something about our minds leads us to make this distinction, even if it is not valid.

The question is this – what could the difference between conscious and unconscious processes be? Do they rely on different networks in the brain? Do some produce qualia and some not? Do some lead to skilled action and some not? Does the hard problem apply to some but not others? And if so, why? Unless we have a viable theory of consciousness, this apparently natural distinction implies what we might call a 'magic difference'.

To explore these questions, we will consider first perception, then action, and finally how perception and action, and conscious and unconscious, may converge in the phenomena of intuition and creativity.

UNCONSCIOUS PERCEPTION

Suppose you are sitting at dinner, chatting with your friends, oblivious to the hum of the microwave in the corner – until it stops. Suddenly you realise that it was humming along all the time. Only in its silence are you conscious of the noise.

This simple, everyday phenomenon seems odd because – like the more extreme cases of agnosia and blindsight we will consider later in the chapter – it suggests perception without consciousness. It suggests that all along, in some unconscious way, you must have been hearing the noise. It challenges the simplistic notion that perception implies or requires consciousness, and that ‘I’ must know what my own brain is perceiving (Merikle, Smilek, and Eastwood, 2001).

The phenomena of unconscious (or implicit, or subliminal) perception have been known about since the very early days of psychology. For example, in the 1880s, Charles Peirce and Joseph Jastrow (1885) studied how well they could discriminate different weights by judging the amount of pressure made on the forefinger or middle finger by the end of the beam of a weighing scale. When two were so closely matched that they had no confidence they could tell them apart, they made themselves guess, and to their surprise did better than chance. This was one of the earliest demonstrations of perception without consciousness. At about the same time, Boris Sidis (1898), another American psychologist and friend of William James, showed volunteers letters or digits on cards so far away they could barely see them, let alone identify them. Yet when he asked them to guess, they also did better than chance. In both cases, people deny consciously detecting something while their behaviour shows that they *have* detected it.

Sidis concluded that his results showed ‘the presence within us of a secondary subwaking self that perceives things which the primary waking self is unable to get at’ (1898, p. 171). As Dan Wegner (2005) points out, the idea of this ‘subliminal self’ implies the existence of its alter ego: a real or conscious self which is capable of fine thoughts and freely willed actions. This is a trap that psychology still falls into, he argues. Whenever there is talk of automatic behaviour, unconscious processes, or subliminal effects, there is an implicit comparison with conscious processes, yet those remain entirely unexplained. He even suggests that ‘psychology’s continued dependence on some version of a conscious self makes it suspect as a science’ (Wegner, 2005, p. 22).

Even if we reject Sidis’s notion of the two selves, his results clearly seemed to demonstrate unconscious perception. However, resistance to this possibility was extraordinarily strong right from the start and continued that way for most of a century (Dixon, 1971; Dijksterhuis, Aarts, and Smith, 2005).

In the early experiments, conscious perception was defined in terms of what people said. This fits with the common intuition that each of us is the final arbiter of what is in our own consciousness: if we say we are conscious or unconscious of something, then (unless we are deliberately lying) we are. Yet this intuition is problematic for several reasons.

One problem is that whether people say they have consciously seen (or heard or felt) something depends on how cautious they are being. This became clearer in the mid-twentieth century with the development of signal detection theory.

‘Even today, when the reality of unconscious perception has been confirmed beyond reasonable doubt [...] there remains almost unshakeable resistance’

(Dixon, 1971, p. 181)

This mathematical theory requires two variables to explain how people detect things like sounds, flashes of light, or touches on the skin. One variable (d' or d -prime) is the person's sensitivity (how good their eyes are, how acute their hearing is). The other, β , is their response criterion (how willing they are to say 'yes, I see it' when they are unsure). These two can vary independently of each other.

Most relevant here is that without their realising it, and with exactly the same sensitivity, people can apply a different criterion. For example, if there is a financial incentive to detect a light flash and no penalty for a false positive, then most people will set a very lax criterion, but if saying 'I see it' when it's not there makes them look stupid or lose money, then they set their criterion much higher.

This means that there is no fixed threshold (or limen) that separates the things that are 'really seen' or 'really experienced' from those that are not. It implies, once again, a difficulty with the idea that things are unequivocally either 'in' consciousness or 'out' of consciousness, and makes the concepts of subliminal and supraliminal perception much more complicated. Some have argued for abandoning the term 'subliminal' altogether in favour of 'implicit' or 'unconscious', but this does not really solve the problem, and generally the term has been retained.

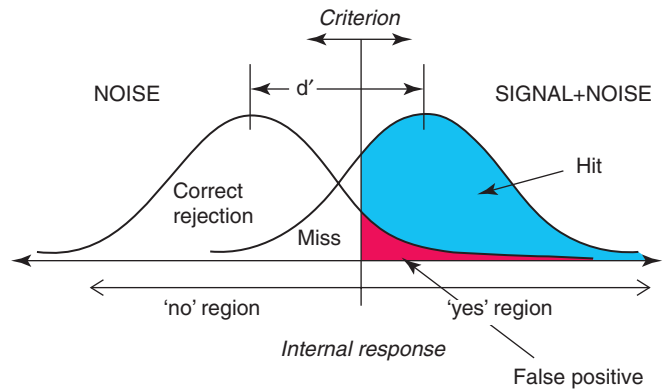


FIGURE 8.1 • Signal detection theory shows how to measure the ability to detect a signal embedded in noise, such as when we try to hear a faint sound or see a briefly flashed word or letter, and replaces the idea of a fixed threshold. The left curve represents just noise and the right curve noise plus signal, and the task is to decide whether there is really a signal or not. A person's criterion may shift even if their sensitivity (d') does not. Areas under the curves show the proportion of hits, misses, correct rejections, and false positives.



PRACTICE 8.1

DID I DO THIS CONSCIOUSLY?

As many times as you can, every day, ask yourself '**Did I do this consciously?**'

You might get out of bed, put on a T-shirt, pick up your toothbrush, or carry out any number of small actions. After any of these, ask the question.

Does asking the question itself make a difference?

Another problem stems from the behaviourist suspicion of verbal reports, with some wanting more reliable, 'objective' measures of consciousness than what people say. But this is a rather curious idea. On the one hand, making a verbal report by speaking or writing is just as much an objective action as is pressing a button or pointing. For this reason, we will not refer to verbal reports as 'subjective measures', as some writers do. But on the other hand, if *all* objective measures of discrimination are taken as evidence of *conscious* perception, then evidence for *unconscious* perception seems to be ruled out by definition (Kihlstrom, 1996). In other words, this move would define away the whole idea that people might be able to demonstrate, by their behaviour, detection of stimuli which they said they were not conscious of.

• SECTION THREE: BODY AND WORLD

Despite the confusion, in the 1970s and '80s these objections prompted progress in both research methods and theory. The basic requirement was to demonstrate a dissociation between two measures: a 'direct' measure, taken to indicate conscious perception, and an 'indirect' measure, to indicate unconscious perception. British psychologist Tony Marcel (1983) adapted the method of semantic priming, in which one word (the prime) influences the response to a second word (the target). For example, if the prime and the target word are semantically related (e.g. doctor and nurse), recognition of the target is faster. Marcel made such primes undetectable by flashing a visual mask immediately after them, yet semantic priming still occurred. This seemed to mean that people's word recognition was affected by primes they did not see: when asked to say whether a word had been presented before the target one, or to guess the masked word, their verbal responses made clear that they were unaware of the words' presence. Other kinds of masked priming were also used, but controversy ensued because although some people successfully replicated the effects, others failed to.

The controversy was resolved when Canadian psychologists James Cheesman and Philip Merikle (1984, 1986) proposed a distinction between what they called the 'objective threshold' and the 'subjective threshold'. The objective threshold is defined as 'the detection level at which perceptual information is actually discriminated at a chance level', whereas the subjective threshold is 'the detection level at which subjects claim not to be able to discriminate perceptual information at better than a chance level' (1984, p. 391). The latter, naturally, is higher than the former.

Cheesman and Merikle used a Stroop-priming task in which participants had to name a colour after being primed with a colour word for different lengths of time. Congruent colour words reduce reaction time, but incongruent words increase it (the Stroop effect). The question was whether primes presented so briefly as to be undetectable would affect reaction times. Cheesman and Merikle measured participants' objective threshold using a reliable four-alternative forced-choice procedure, and their subjective threshold by asking them to judge their own ability to discriminate the words. They found a priming effect (i.e. evidence for unconscious perception) when the length of time between the prime and the mask was below the subjective threshold, but none at all when it was below the objective threshold.

Their conclusion was that unconscious perception occurs primarily when information is presented below the subjective threshold but above the objective threshold. They were then able to show that previous experiments had confused the two, with some measuring one and some the other. From this we can conclude that the objective threshold really is the level below which stimuli have no effect of any kind, but there is a level above that at which a stimulus can have an effect even though the person denies being conscious of it.

Many of these experiments used words as stimuli and implied the possibility of unconscious semantic analysis. This possibility has been debated for a century or more, with many arguments about just how much meaning someone can extract from a stimulus they deny seeing. Psychologist John Kihlstrom concludes that 'With respect to subliminal stimulation, the general rule seems to be that

the further the stimulus moves from [i.e. below] the subjective threshold, the less likely it is to be subject to semantic analysis' (1996, pp. 38–39). This is just one example showing that the effects of stimuli can differ when they are above or below different thresholds.

This idea of two thresholds creates difficulties if we think that all processes must be either conscious or unconscious. We can also go one step further and suggest that if human behaviour is controlled by multiple parallel systems without an inner controller or central self in charge, then we might expect to find different systems responding at many different thresholds, rather than just two. So, when thinking about implicit, unconscious, or subliminal processes, we must remember that there may be many thresholds, none of them fixed or unvarying.

The fascination and fear of the subliminal took a popular turn in 1957 when James Vicary claimed to have massively increased sales of Coke and popcorn by flashing the messages 'Drink Coca-Cola' or 'Hungry? Eat Popcorn' very briefly during a film. Vicary's study seems to have been a publicity hoax, but many people still fear the power of subliminal advertising. And subliminal messages may have small effects. For example, priming thirsty participants with brand names can affect their intention to drink a particular brand of drink (Karremans, Stroebe, and Claus, 2006). However, the effects on actual behaviour seem to be weak (Dijksterhuis, Aarts, and Smith, 2005).

'subliminal self-help messages do not have any effect at all'

(Dijksterhuis, Aarts, and Smith, 2005, p. 77)

Also popular is the idea that subliminal self-help programmes can reduce anxiety, improve self-esteem, health, and memory, or help people to give up smoking or lose weight. Software designed to insert subliminal messages while you sleep, work, or play games is also commercially available. There is evidence of a placebo effect that depends on what the label says rather than what is inside; this may explain why people keep buying them, but there is no evidence of an effect of the messages themselves. Merikle concludes that 'There is simply no evidence that regular listening to subliminal audio self-help tapes or regular viewing of subliminal video self-help tapes is an effective method for overcoming problems or improving skills' (2000, p. 499).

The question about thresholds of awareness becomes crucial in the case of anaesthesia. In ordinary medical practice, one in every thousand to two thousand patients has general anaesthetic wrongly administered or monitored, and experiences some awareness during anaesthesia. In the worst cases, patients experience severe pain and fear but can do nothing to make their state known. This is somewhat like locked-in syndrome but is at least temporary. The very possibility of 'unintended intra-operative awareness' was long denied, but improved understanding of the four separable functions of anaesthetic – paralysis, pain-numbing, amnesia, and loss of consciousness – mean that it is now generally accepted, and great efforts are made to prevent it happening.

Being unresponsive and amnesic under anaesthetic is not necessarily the same as being rendered unconscious (Alkire, Hudetz, and Tononi, 2008). A somewhat unnerving effect was found in a controversial study of people undergoing general anaesthesia (Levinson, 1965). A mock crisis was staged during a real operation by the experimenter reading out a statement to the effect that the patient

was going blue and needed more oxygen. A month later the ten patients were hypnotised and asked whether they remembered anything that had occurred during their operation. Four of the ten remembered the statement almost verbatim, and a further four remembered something of what was said. This conjures up visions of people being unconsciously affected by horrific scenes from operating theatres.

Generally, however, unconscious processing during full anaesthesia can be detected, but the effects are small. For example, explicit memories for information presented under anaesthetic may be retrieved only if testing occurs within thirty-six hours, and have little or no effect on postoperative recovery, while priming effects may depend on the specific anaesthetic used (Merikle and Daneman, 1996; Kihlstrom and Cork, 2007).

In one fascinating technique, a tourniquet is applied to a patient's forearm before the anaesthetic. This means that the hand is not paralysed and patients can sometimes have a conversation using hand signals, although afterwards they deny ever having been awake: 'Thus, retrospective oblivion is no proof of unconsciousness' (Alkire, Hudetz, and Tononi, 2008, p. 877). It is tempting to believe in a pivotal point somewhere between behavioural unresponsiveness and a flat EEG (one of the criteria for brain death) where 'consciousness must vanish'. But EEG indexing sometimes still yields the wrong result, for example in cases using the isolated forearm technique. 'Either the EEG is not sensitive enough to the neural processes underlying consciousness, or we still do not yet fully understand what to look for' (p. 877).

'Either the EEG is not sensitive enough to the neural processes underlying consciousness, or we still do not yet fully understand what to look for'

(Alkire, Hudetz, and Tononi, 2008, p. 877)

EMOTIONAL AND SOCIAL EFFECTS

Some of the most striking experiments on unconscious perception concern the emotional effects. It is well known that people prefer familiar things – including simple images they have seen before. This is the 'mere-exposure effect' and, perhaps surprisingly, it works for subliminal stimuli, too. In a famous study (Kunst-Wilson and Zajonc, 1980), participants were shown meaningless shapes so briefly that none reported seeing them. Responses were measured in two ways. In a recognition task, participants had to choose which of two shapes had been presented before. They could not do this and scored at chance (50%). Next they were asked which of the two they preferred, and this time chose the one they had seen before 60% of the time. This is a good example of how two different objective measures of awareness, neither of them a direct verbal report about the 'contents of consciousness', can lead to different answers.

It is tempting to ask which measure reveals what the participants were *really* conscious of. This question is a natural one for the Cartesian materialist, who believes that there must be an answer: things must be either in or out of consciousness. An alternative is to reject this distinction and say that there is no ultimately 'correct' measure of whether someone is conscious of something or not; there are just different processes, different responses resulting from or accompanying those processes, and different ways of measuring them. In this view, there is no answer to the question 'what was I really conscious of?'

In a later experiment, participants had to rate a series of unfamiliar Chinese ideographs according to whether they thought each represented a 'good' or a 'bad'

concept (Murphy and Zajonc, 1993). One group of participants saw either a smiling face or a scowling face for one second before each ideograph. They were told to ignore these faces and concentrate only on rating the ideographs. The second group were shown the faces for only 4 msec, which is not long enough to see them. The striking result was that the first group managed to ignore the faces, but the second group were influenced by the faces they claimed not to see. If the invisible face was smiling they were more likely to rate the ideograph as 'good'.

Outside the lab, such effects may permeate our complex social worlds as we unconsciously imitate other people's facial expressions, mannerisms, moods, and tone of voice, or make spontaneous judgements about people without knowing why. Many of these judgements depend on emotional signals that are difficult to hide, and are interpreted remarkably fast and accurately (Choi, Gray, and Ambady, 2005). Other 'implicit impressions', however, are less reliable. For example, if photos of people with a distinctive physical feature are paired with positive or negative events, this affects subsequent responses to other people who have that same physical feature. You may treat strangers more positively if they resemble someone you love. There is also the phenomenon of spontaneous trait transference, in which descriptions of a person's traits are transferred to the person who gave the description. For example, if you describe someone as kind and clever, or cruel and devious, the listener may unconsciously attribute those characteristics to you (Uleman, Blader, and Todorov, 2005), implying a 'boomerang effect' of malicious gossip.

What is going on in the brain during unconscious or implicit perception? Studies have shown that both positive and negative faces can produce significant changes in amygdala activation, even when the stimuli are not consciously perceived (Williams et al., 2004). Using dichoptic colour fusion, Moutoussis and Zeki (2002) made binocularly viewed face and house stimuli invisible and showed that the relevant brain areas (the fusiform face area and parahippocampal place area, respectively) were still activated, if less strongly, even when the stimuli were not perceived. This finding could be interpreted in two quite different ways. Perhaps at a certain level of activity the stimulus becomes conscious, or alternatively at that level the activity starts to have other effects within the brain and body.

Stanislas Dehaene and his colleagues used ERPs (event-related potentials) and fMRI to investigate how the presentation of numbers masked and presented too briefly to be seen can help in subsequent processing of related numbers, indicated by the speed of pressing a response key. They found activity in motor as well as sensory areas, suggesting covert responses to the primes that could not be reliably reported or discriminated. They concluded that 'A stream of perceptual, semantic, and motor processes can therefore occur without awareness' (Dehaene et al., 1998, p. 597). This is reminiscent of William James's contention that 'Every impression which impinges on the incoming nerves produces some discharge down the outgoing ones, whether we be aware of it or not' (1890, ii, p. 372).

This brings us back to that troublesome 'magic difference' and the hunt for the NCCs. If every impression produces effects in the brain, then what is the difference between the ones we are aware of and those we are not? In the following section, we will face a similar question as it applies to the difference between conscious and unconscious action.

'there is still no agreement as to the role of unconscious or preconscious cognitive processes'

(Merikle, 2007, p. 512)

CONSCIOUS AND UNCONSCIOUS ACTIONS

There is no doubt that we seem to do some things consciously, others unconsciously, and others sometimes one way and sometimes the other. On this basis, we can divide actions into what our intuition suggests are five distinct types.

- 1 Always unconscious. I can consciously wiggle my toes or sing a song, but I cannot consciously grow my hair or increase my blood sugar level. Spinal reflexes that depend on neural connections outside of the brain are always unconscious, and much of visuomotor control is carried out too fast for consciousness to play a role.
- 2 Usually unconscious. Some actions that are normally carried out unconsciously can be brought under conscious control by giving feedback about their effects, or 'biofeedback'. For example, if a visual or auditory display is provided to indicate when your heart beats faster or slower, when your left hand is warmer than your right, or when your palms sweat more, you can learn to control these variables, even when obvious actions that might produce the changes, such as clenching your hands or jumping up and down, are prevented. The sensation is rather odd. You know you can do it, and feel in control, but you have no idea *how* you do it. This should remind us that the same is true of most of what we do. We may consciously open the door but have no idea how all the intricate muscular activity required to turn the handle is coordinated. The whole action seems to be done consciously while the details remain unconscious.
- 3 Initially conscious. Many skilled actions are initially learned with much conscious effort but with practice come easily and smoothly. While biofeedback moves actions into conscious control, automatization does the reverse. You probably first learned to ride a bicycle with the utmost conscious concentration. Learning any motor skill is like this, whether it is skateboarding or skiing, using a mouse or keyboard, or learning the movements in yoga or Tai Chi. After complete automatization, paying conscious attention can even be counter-productive, making you fall off your bike or struggle to even walk normally.
- 4 Either conscious or unconscious. Many skilled actions, once well learned, can be done either way. Sometimes Sue makes a cup of tea with utmost mindfulness, but often she finds she has put the kettle on, warmed the teapot, found the milk, made the tea, and carried it back to her study without, apparently, being conscious of any of the actions. As we learnt in the previous chapter, changing the quality of attention through mindfulness is a powerful way of deautomatising actions and perceptions. The classic example is the unconscious driving phenomenon (see Chapter 5). Here we have detailed, complex, and potentially life-threatening decisions being made correctly without, apparently, any conscious awareness.
- 5 Always conscious? Finally, some actions seem always to be done consciously. For example, when I try to remember a forgotten name or a route in an unfamiliar city, I seem to struggle consciously, while a familiar name trips effortlessly off the tongue. When I have to make a difficult moral decision or am composing a poem, I seem to be far more conscious than when deciding what clothes to put on. It is tempting to say that these kinds of thinking, decision-making, or creativity *require* consciousness.

Here, then, is the rub. If the same action is carried out on one occasion consciously and on another occasion unconsciously, what is the difference? Obviously, there is a phenomenal difference – they *feel* different – but why?

We must avoid jumping to unwarranted conclusions. For example, we might start by observing that we made a difficult moral decision *consciously* while we made the tea *unconsciously*; jump from there to the conclusion that the former *requires* consciousness, while the latter does not; and finally, to the conclusion that consciousness itself *does* the deciding. But this is not the only interpretation. Another possibility is that the processes involved in making difficult moral decisions incidentally give rise to the impression of their being done consciously, while those involved in making tea do not. Another is that difficult tasks require more of the brain to be involved or more parts to be interconnected, and this greater connectivity either is, or gives rise to, the phenomenal sense of doing the action consciously. Whenever we compare actions done with and without consciousness, we must remember these different interpretations. This is relevant to perception (Chapter 3), the neural correlates of conscious and unconscious processes (Chapter 4), the Cartesian theatre (Chapter 5), intuition and unconscious processing (later in this chapter), and the nature of free will (Chapter 9), but for now the question concerns the role of consciousness in action – what is the difference between actions performed consciously and those done unconsciously?

If you believe that consciousness has causal efficacy (i.e. *does* things), then you will probably answer that consciousness caused the former actions but not the latter. (See the website for an additional Activity on this.) In this case, you must explain how subjective experiences can cause physical events. If you do not think that consciousness can *do* anything, then you must explain the obvious difference some other way. Theories of consciousness differ considerably in their answers, as we can see from the following examples.

**DID I DO THIS
CONSCIOUSLY?**

THEORIES

CAUSAL THEORIES

Some theories have a clear causal role for consciousness, most obviously dualist theories, but they face the problem that for consciousness to have any effects it must interact with matter. Descartes located this interaction in the pineal gland, but he could not explain how it worked. Two centuries after Descartes, in his *Principles of Mental Physiology*, William Benjamin Carpenter (1874) proposed that in one direction physiological activity excites sensational consciousness, while in the other direction sensations, emotions, and volitions liberate the nerve-force with which the appropriate part of the brain is charged. But he too could not explain how.

A century later, Popper and Eccles's (1977) dualist interactionism faced exactly the same problem in having to explain how the independent 'self-conscious mind' could interact with the 'liaison areas of the dominant cerebral hemisphere' (p. 362). Eccles later proposed (1994) that all mental events and experiences are composed of 'psychons' and every psychon interacts with one dendron in the brain. Although this localises the interaction, he could not explain how it worked.

• SECTION THREE: BODY AND WORLD

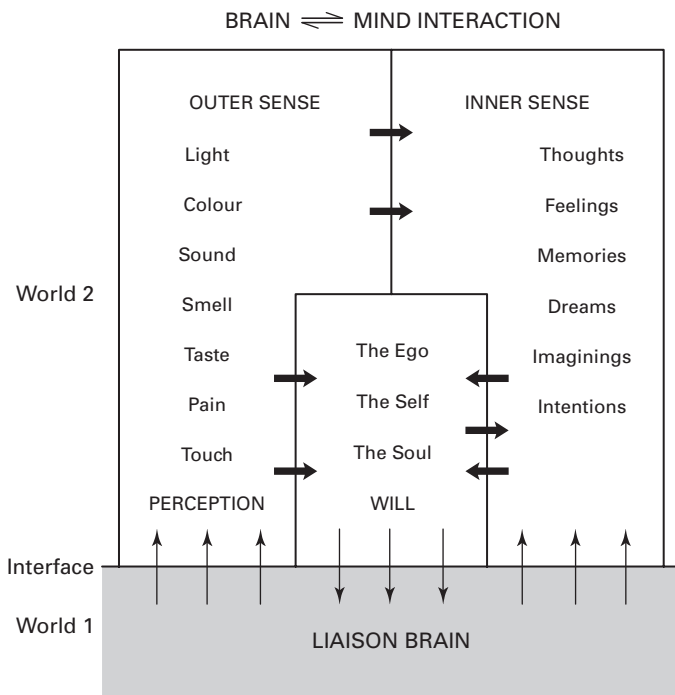


FIGURE 8.2 • How the brain interacts with the mind, according to Popper and Eccles. The three components of World 2 (mental) are outer sense, inner sense, and the ego or self, shown here with their connectivities. Also shown are the lines of communication between World 1 (physical) and World 2, i.e. from the liaison brain to and from these World 2 components. The area of the liaison brain has a columnar arrangement, and is supposed to be enormous, including a hundred thousand or more open modules (Popper and Eccles, 1977, p. 360).

Libet's conscious mental field (CMF) also acts both ways, providing in one direction 'the mediator between the physical activities of nerve cells and the emergence of subjective experience', and in the other 'a causal ability to affect or alter some neuronal functions' (Libet, 2004, p. 168).

These theories unambiguously answer our question. When an action is carried out consciously, the self-conscious mind or the CMF causes the brain to carry it out by interacting with it; when an action is carried out unconsciously, the brain acts alone. But the interactions remain unexplained.

British psychologist Jeffrey Gray calls dualism 'moribund' (2004, p. 73) but also tries to retain a causal role for 'consciousness in its own right' (p. 90). He proposes that 'The decision [to act] is made by the unconscious brain and enters conscious awareness only after the event' (p. 92). Consciousness itself monitors features relevant to ongoing motor programs, and permits change in the variables controlled by unconscious servo-mechanisms, or feedback-based controllers. For

Gray, qualia are 'raw feels' created by the brain which, once created, can be put to use in a great variety of cognitive processes. How either conversion process works is unexplained, but the answer to our question would presumably be that only in conscious actions are qualia put to use.

Finally, there is global workspace theory (Chapter 5). According to Bernard Baars, consciousness is a supremely functional biological adaptation. It is a kind of gateway: 'a facility for *accessing, disseminating, and exchanging information*, and for *exercising global coordination and control*' (1997b, p. 7). The nine functions of consciousness include being 'essential in integrating perception, thought, and action, in adapting to novel circumstances, and in providing information to a self-system' (p. x). He firmly rejects the idea 'that consciousness simply has no causal role to play in the nervous system' (p. 165).

Baars illustrates the power of consciousness with the following example. Imagine that as you are reading this book, you become aware of a strange, foetid animal smell, the noise of heavy hooves, and hot breath down the back of your neck. Although reluctant to stop reading, you suddenly have the wild thought that there might be a large animal in the room. You turn your head, see the large angry ferocious bull, and leap from your chair. Consciousness, at least in our evolutionary past, would have saved us from danger, he says. The problem with this interpretation lies in the timing. The results of many experiments suggest that you would have leapt out of that chair long before you could have consciously thought about the danger (Chapter 9).

On Baars's theory, the answer to our question is that actions which are performed consciously are shaped by conscious feedback while unconscious actions are not. For example, you might unconsciously make a speech error, but when you hear

the mistake you can put it right because consciousness creates global access to further unconscious resources. Yet it is not entirely clear whether consciousness is supposed to be the cause, the same as, or the result of access to the GW (Rose, 2006).

Similar ideas appear in Dehaene's neuronal global workspace theory, in which consciousness has clear causal power. As he puts it, 'consciousness has a precise role to play in the computational economy of the brain – it selects, amplifies, and propagates relevant thoughts' (2014, p. 14). More specifically, for example, it provides a summary of the environment to help guide action (p. 100). For Dehaene, anything that we are aware of, because it has reached the conscious workspace, 'becomes available to drive our decisions and our intentional actions, giving rise to the feeling that they are "under control"' (2014, p. 167). Again, there are two ways of interpreting this: one is that the contents of the GW somehow 'become conscious' and this gives them causal power; the other is that the contents have effects simply because they are broadcast from the GW to many other brain areas. One involves an unexplained transformation from unconscious to conscious and the other does not. But either way, the answer to our question is that conscious actions have access to the global workspace, while unconscious actions do not.

'consciousness has a precise role to play in the computational economy of the brain'

(Dehaene, 2014, p. 14)

NON-CAUSAL THEORIES

At the other extreme are theories which reject the idea that consciousness can *cause* events (see Chapter 1). One example is eliminative materialism, which denies the existence of consciousness as anything distinct from its material basis. Epiphenomenalism accepts the existence of consciousness but denies that it has any effects. In its traditional form, this is a somewhat strange idea, implying a causal chain of events leading from sensory input to behaviour, with consciousness produced as a by-product that has no further effects at all. As we have seen, one apparent stumbling block here is that if consciousness had no effects we could not even talk about it, let alone write a book about it. Epiphenomenalism in this form is highly counter-intuitive, and it cuts against some of our most dearly cherished intuitions, 'entailing that what we believe, feel, sense, remember, etc., does not make a causal difference to what we do' (Pauen, Staudacher, and Walter, 2006).

In the philosophy of mind, there are two main representational theories: 'higher-order perception' (HOP) theory and various types of 'higher-order thought' (HOT) theory (Carruthers, 2007). Other higher-order theories include HOGS (higher-order global states) and HOST (higher-order syntactic thought) theories (Gennaro, 2004, 2017). According to HOP theory, being conscious of a mental state means monitoring first-order mental states in a quasi-perceptual way – with something analogous to an 'inner eye' or 'inner sense' (Lycan, 2004). According to HOT theories, a mental state is conscious if the person has a higher-order thought to the effect that they are in that state (Rosenthal, 1995, 2008). For example, my perception of a red flash is conscious only if accompanied by a HOT that 'I am seeing a red flash'.

Higher-order theories readily answer our questions. What is the difference between actions performed consciously and those done unconsciously? Answer – there are HOPs or HOTs about them. No special place or kind of neuron is required; only that the brain must construct HOPs/HOTs. Although HOPs/HOTs have effects (i.e.

'Animals have plenty of access to their experiences, but probably little in the way of higher order thought about them'

(Block, 2005, p. 50)

making things like actions conscious), they do not cause actions but are more like a commentary on them. Indeed, an HOP/HOT may take time to construct and so may happen *after* an action that is experienced as performed consciously, which seems to fit with evidence discussed below and in Chapter 9. On these theories, zombies, though conceivable, are not possible, because anything behaviourally identical to us (in being able to report HOPs/HOTs) would by definition be conscious.

However, such theories face difficulties, such as deciding what counts as the content of an HOP/HOT (e.g. what sort of ideas about redness are involved in the thought?). They also mean denying consciousness to creatures incapable of HOTs, and have trouble dealing with states which seem to be conscious without thought or an observer of any kind (Seager, 2016), especially mystical experiences and deep states of meditation (Blackmore, 2011).

In the end, though, we return to our familiar question: why should the posited extra ingredient (here, targeting a mental state with a HOP/HOT) cause that state to be conscious? Can we imagine zombies with lots of higher-order thoughts but no consciousness? Or is having an HOT about your experience like asking 'am I conscious now?' and always getting the answer yes, like the fridge light that is always on when you open the door to look?

FUNCTIONALISM

Functionalism, like so many other words to do with consciousness, is used in many different and sometimes contradictory ways. Within the philosophy of mind, it is the view that mental states are functional states. So, for example, if someone is in pain, the pain is understood in terms of the input from the damage done, the output of behaviours such as crying or rubbing the wound, and other mental states such as the desire for the pain to go away, which can also be specified functionally. This means that any system which executed exactly the same functions as a human being in pain would also be in pain, so zombies are impossible. Functionalism is often opposed to physicalism because it emphasises the functions a system carries out rather than what it is physically made of, and to behaviourism because it considers internal functions and not only behaviour. But the implications for subjective experience are not obvious. A common view is that functionalism works well for explaining some mental states but is much less clear in accounting for phenomenal consciousness or qualia (Van Gulick, 2007) – but note that philosophers take 'mental states' to include such things as desires and beliefs, which other disciplines do not.

The term is also used, especially by psychologists and in discussions of artificial intelligence (Chapter 12), to mean that any system which could carry out exactly the same functions as a conscious system would also, necessarily, be conscious. This is the idea of multiple realisability: the same conscious state could be realised in multiple ways as long as the same functions were carried out. If we ask what the difference is between actions carried out consciously and those carried out unconsciously, the functionalist will answer in terms of the different functions involved; there is no separate consciousness to play a causal role. Although many of the theories in this book are, broadly speaking, functionalist, including representationalist theories such as HOP and HOT theories (Kobes, 2007), they struggle to explain how or why functions can *be* phenomenal consciousness, and to explain (or explain away) qualia.

We started with the simple idea that consciousness causes at least some of our actions, but the theories and experiments discussed reveal serious problems with this commonsense notion. So a gentle alarm bell should ring every time we read that consciousness directs attention, or gives us the ability to introspect; that it drives our emotions and our higher feelings; or that it helps us assign priorities or retrieve long-term memories. Comments such as this are deeply embedded in our ordinary language about consciousness, and can easily be found in the writings of psychologists, philosophers of mind, and others. It is not obvious which, if any, of them is true. Maybe exploring an everyday example of skilled action will help us decide.

THE CAUSAL EFFICACY OF CONSCIOUSNESS

Imagine someone throws you a ball and you catch it – or to make things more realistic, scrumple a piece of paper up into a ball right now, throw it up in the air, and catch it yourself. Do this a few times and ask yourself what the role of consciousness was in this simple skilled action. You were conscious of doing the catching, and of the sight of the ball as your hands reached for it, but did the consciousness itself cause anything to happen? Without consciously seeing the ball, could you have caught it?

In doing this simple task, the causal sequence *seems* to be 1) consciously perceive, and 2) act on the basis of the conscious experience. This is sometimes known as the ‘assumption of experience-based control’.

When you think about it, this is a strange notion. It means two mysterious conversions: first the physical information in nerve firings in the visual system must somehow be turned into conscious experiences, and then the conscious experience must somehow act back on the brain, causing more nerve firings to direct the appropriate action. But if consciousness is subjectivity (experience, non-physical qualia, what it’s like to be), how can either process work? How can non-physical experiences *cause* physical firings of nerve cells or movements of muscles? And where and how does this consciousness bit happen in the brain?

This is one example of Susan Hurley’s ‘Classical Sandwich’ model of the mind: ‘The mind is a kind of sandwich, and cognition is the filling’ (2001, p. 3). You start with the bread of perception, then you have the filling of cognition, and on top you have the bread of action. Now we are right back to the mind–body debate and equally ancient discussions about the problem of mental causation. In ancient Greece, philosophers were already debating how mental states such as beliefs, desires, or thoughts could have physical effects. Descartes faced the problem when he began to think about the human body as a mechanism and realised that emotions and volition did not easily fit in. He took a dualist way out – which, as we have seen, almost certainly does not solve the problem.

Note that although the debate about mental causation goes back centuries, the question of consciousness is just one part of it. As philosopher Jaegwon Kim (2007) points out, even if a conscious thought can cause other thoughts or actions, this doesn’t necessarily mean that the fact that the thought was conscious, rather than (say) its content was the relevant factor. What we want to know here is whether consciousness has any causal power.



FIGURE 8.3 • Descartes tried to explain reflex responses, like removing your foot from a hot fire, in purely mechanical terms. He believed that the fire affected the skin and pulled a tiny thread which opened a pore in the brain's ventricle and caused animal spirits to flow. But where does consciousness come in? It is tempting to think that a signal must come 'into consciousness' before we can decide to act on it. But is this right?

'to urge the automaton-theory upon us [...] is an unwarrantable impertinence in the present state of psychology'

(James, 1890, i, p. 138)

'from a third-person perspective, phenomenal consciousness appears to play no causal role in mental life, while from a first-person perspective it appears to be central'

(Velmans, 2009, p. 315)

In the nineteenth century, as physiologists began to understand reflex arcs and nerve function, this problem loomed even larger. Shadworth Hodgson declared that feelings, however intensely they may be felt, can have no causal efficacy whatever. He likened the states of the nervous system to the stones of a mosaic, and feelings, or 'states of consciousness', to the colours on the stones (1870, i, p. 336). All the work of holding the mosaic in place is done by the stones, not by the colours. In other words, conscious states are epiphenomena. This was similar to Thomas Huxley's claim that we humans are 'conscious automata' (Chapter 1). James, however, objected that 'to urge the automaton-theory upon us [...] is an unwarrantable impertinence in the present state of psychology' (1890, i, p. 138). But his reasons for saying so were, as much as anything, about acknowledging our ignorance and appealing to the idea of common sense – which we have already seen can be misleading.

James predicted that for years to come, we would have to infer what happens in the brain by making inferences from our feelings or behaviours: 'The organ will be for us a sort

of vat in which feelings and motions somehow go on stewing together, and in which innumerable things happen of which we catch but the statistical result' (p. 138). More than a century later, brain imaging and other technologies have given us far greater knowledge of the stew pot or the sandwich filling. But the conundrum is far from resolved, and may even have become worse, perhaps in part because of an overly tight focus on the brain in isolation. We may think that our subjective feelings and conscious volitions cause our actions, yet when we study the intricate workings of the brain there is no room for them to do anything at all. Information enters the nervous system through the senses, flows through numerous parallel pathways to various brain areas, and ultimately affects a person's speech and other actions. But where do the conscious sensations and volitions come in? How could they intervene – or why should they – in such a continuous physical process? As Kim (2007, p. 407) puts it: 'Aren't the underlying physical/neural processes ultimately doing all the actual pushing and pulling, with no work left for consciousness to do?'

For Max Velmans, 'consciousness presents a Causal Paradox'. As he points out: 'Viewed from a first-person perspective, consciousness appears to be necessary for most forms of complex or novel processing. But, viewed from a third-person perspective, consciousness does not appear to be necessary for any form of processing' (2009, p. 300). Taking an example from medicine, he notes that we take all four possible causal links between physical and mental for granted (biomedical interventions, neurosurgery and psychoactive drugs, psychotherapy, and psychosomatic medicine). An adequate theory of consciousness, he says, must make sense of these causal interactions and so resolve the paradox without violating either our intuitions about our own experiences or the findings of science.

THE ROLE OF CONSCIOUSNESS IN SKILLED ACTION

The fastest serves cross the tennis court in a little over half a second, and the ball starts its flight at well over a hundred miles an hour. Yet these super-fast serves can be returned with stunning accuracy. Does the receiver have time for the mysterious double conversion to consciousness and back? Is conscious perception even necessary for such skilled movements?

The answer is no. Studies of skilled motor actions reveal a dissociation between fast visuomotor control and conscious perception. For example, in some experiments participants are asked to point at a visual target; then just as they begin to point, the target is displaced. If the displacement is made during a voluntary saccade, participants do not notice the displacement even though they rapidly adjust their arm movement to point correctly at the final position (Bridgeman et al., 1979; Goodale, Pelisson, and Prablanc, 1986). In other words, their behaviour is accurately guided by vision even though they do not consciously see the target move. Accurate movements can also be made towards stimuli that are not consciously perceived at all. When small visual targets were made invisible by presenting a larger stimulus 50 ms later (this is called backward masking), participants still responded correctly to the target they claimed not to have seen (Taylor and McCloskey, 1990).

In the case of the tennis serve, or catching your scrumpled paper ball, the ball *is* consciously perceived – but when? Does the conscious perception occur soon enough to affect the action?

One experiment (Paulignan et al., 1990) asked participants to track by hand a moving object that was suddenly displaced. They were able to respond within about 100 ms, but when asked afterwards to estimate at which point in the movement they had seen the displacement, they consistently reported that the object jumped just when they were about to touch it – that is, much later than either the actual displacement or their own corrective movement. This finding suggested that conscious awareness may come too late to play a causal role in the action.



ACTIVITY 8.1

Incubation

Incubation is the process of putting a problem ‘on the back burner’, just allowing a solution to come by itself – if it will. Three steps are required. First, you have to do the hard work of struggling with the problem or acquiring the necessary skills. Second, you have to drop the struggle and leave the problem to itself, perhaps by engaging in some other activity, or just sleeping on it. In this second stage, any conscious effort is likely to be counter-productive. Third, you have to recognise the solution when it appears.



FIGURE 8.4 • Move three coins to turn the triangle upside down.




FIGURE 8.5 • Move two match sticks to make the cow face the other way. You can try this one on your friends using real match sticks; leave it on the bar or the dinner table and let them incubate it too.

Here are three simple brain-teasers to practise incubation. If you are working on your own, have a good go at trying to solve them, until you get really frustrated. Then forget all about them and read more of the book, or do something else for half an hour or so. When you come back to the problem, you may find that the solution just ‘pops into your mind’. If you are working in a group, you can start a lecture or discussion with five minutes working on the problems and then return to them at the end, making sure that those people who solve the problems quickly don’t give the answers away and spoil the experience for everyone else. The solutions are given in Figures 8.17–8.19 on p. 216.

• SECTION THREE: BODY AND WORLD




FIGURE 8.6 • Draw just two squares to provide each pig with its own enclosure.

Another study soon afterwards (Castiello, Paulignan, and Jeannerod, 1991) found out more by timing both motor responses and subjective awareness in the same experiment. Participants sat at a table facing three translucent dowels, any of which could light up. Their task was to watch for a light, shout 'Tah' as soon as they saw it, and grab the lit dowel as fast as they could. In the first session, a single dowel was lit many times and motor and vocal reaction times were measured. In a second session of 100 trials, the central dowel was always lit first, but then 20% of the time the light switched to a different dowel as soon as their hand began to move, so that they had to correct their movement.

Reaction time to the initial movement was always about 300 ms, and then in trials where the light switched, a movement correction occurred about 100–110 ms later. On these trials, the participants shouted twice: when the central dowel lit up, the vocal reaction time was 375 ms (the same as in control trials); when the light moved, it was about 420 ms. In other words, the movement was corrected long before the shout that meant 'I've seen it!' The authors argued that the vocal response indicated when the participants became conscious of the light, and concluded that 'neural activity must be processed during a significant and quantifiable amount of time before it can give rise to conscious experience' (p. 2639). As one of them later put it, 'our consciousness has to play catch up' (Jeannerod, in Gallagher, 2008, p. 244).

There are problems with this conclusion. Fast reaction times can be obtained without awareness, so the shout might have been initiated before the participant consciously saw the light move, in which case consciousness might have come even later than estimated. Alternatively it might have come earlier and the full 420 ms have been needed to produce the response. We cannot know because this method does not allow us to precisely time the 'moment of awareness'.

Perhaps, as we suggested in Chapter 5, we should be even more critical and question the very notion of there being a 'moment of awareness' or a time at which the light comes 'into consciousness', because this implies a mental world in which conscious events happen alongside the physical world of brain events. We will return to the question of how to time conscious awareness in Chapter 9. Despite these doubts about timing, the results suggest a dissociation between fast motor reactions and conscious perception. One explanation is that the two are based on entirely different systems in the brain.



FIGURE 8.7 • Venus Williams serves at speeds of up to 125 miles per hour, yet opponents manage to respond. Is there time for a visual signal to 'enter consciousness', be experienced, and then cause a conscious response? Or is this natural way of thinking about our actions misguided?

UNCONSCIOUS ACTION VERSUS CONSCIOUS PERCEPTION?

Milner and Goodale (1995; Goodale and Milner, 2013) suggest a functional dissociation between two vision systems and map this onto the two neural streams in the visual system: the dorsal and ventral streams (Chapter 6). These two streams had often been described as being concerned with spatial vision and object vision respectively, or with the 'where' and 'what' of vision (Ungerleider and Mishkin, 1982). Instead, Milner and Goodale argue for a distinction based on two fundamentally different tasks that the brain has to carry out. One is fast visuomotor control, which needs egocentric models, relating the self to objects in the world; the other is the less urgent visual perception, which needs more allocentric processing – relating objects in the environment to each other. They call these the *vision-for-action* and the *vision-for-perception* systems (Goodale, 2007).

Much of Milner and Goodale's evidence comes from patients with brain damage. One patient, D.F., was taking a shower and was nearly asphyxiated by carbon monoxide poisoning from a faulty water heater. Her partner found her before she died, and when she emerged from her coma, it became clear that her brain had been badly damaged from lack of oxygen. In particular, she was left with visual form agnosia. This means she is unable to recognise the shapes of objects by sight, even though her low-level vision of basic visual features including pattern and colour appears to be intact. She cannot name simple line drawings or recognise letters and digits, nor can she copy them, even though she can produce letters correctly from dictation and can recognise objects by touch. She can, however, reach out and grasp everyday objects (objects that she cannot recognise) with remarkable accuracy.

One experiment with D.F. reveals this extraordinary split between motor performance and awareness. She was shown a vertically mounted disc in which a slot was randomly cut at 0, 45, 90, or 135 degrees. When asked to draw the orientation of the slot, or adjust a comparison slot to the same angle, she was quite unable to do so. However, when given




FIGURE 8.8 • The layout in the experiment by Castiello et al. (1991), showing the arm trajectories both when the lit dowel stays the same and when it changes.



PROFILE 8.1 Melvyn Goodale (b. 1943)



Having emigrated with his parents from England to Canada as a child, Mel Goodale studied psychology before setting off to 'find himself' travelling around Europe. Getting sick of casual jobs and damp apartments and with still no idea what he wanted to do, he headed for graduate school in Calgary, ended up in a lab studying visual neuroscience, and was instantly captivated. He is now Director of the Brain and Mind Institute at the University of Western Ontario, and Co-Director of the Azrieli Program on Brain, Mind, and Consciousness at the Canadian Institute for Advanced Research. He is best known for his work with David Milner on the functional organisation of the visual pathways in the cerebral cortex. Their studies of visuomotor control in neurological patients led to their characterising the two streams of the primate visual system as 'vision for perception' and 'vision for action'. He has since also begun to explore how the blind use echolocation to navigate.




FIGURE 8.9 • The ventral and dorsal visual streams. Ungerleider and Mishkin called them the ‘what’ and ‘where’ streams. Milner and Goodale suggest that they carry out vision for perception and vision for action (or visuomotor control) respectively (Milner and Goodale, 1995, p. 22).

*‘what we think we “see”
is not what guides our
actions’*

(Milner and Goodale, 1995,
p. 177)

a piece of card she could quickly and accurately post it through the slot – a task requiring accurate alignment. How can this be? How can she be unaware of the angle of the slot and yet able to post the card into it? The answer, according to Milner and Goodale, is that she has lost much of the ventral stream that leads to visual perception but retains much of the dorsal stream needed for accurate visuomotor control – losses which have been confirmed by multiple scans (Whitwell, Milner, and Goodale, 2014).

According to Milner and Goodale, these experiments show that sometimes ‘what we think we “see” is not what guides our actions’ (1995, p. 177) and that ‘the visual

signals that give us our experience of objects and events in the world are *not* the same ones that control our actions’ (Goodale, 2014). These findings were subsequently challenged (Franz et al., 2000), reanalysed to meet the challenge (Danckert et al., 2002), and much debated, with an alternative proposal that visual illusions affect the planning of actions but not their online control (Glover, 2002; Goodale, 2007).

These dissociations are not limited only to brain-damaged patients. The same separation between perception and motor control was reported in a study using visually normal participants tricked by a visual illusion (Aglioti, Goodale, and DeSouza, 1995). Thin discs were made to look different sizes by surrounding them with rings of larger or smaller circles. This is the Ebbinghaus or Titchener illusion, which has been shown to trick even fish (Sovrano, Albertazzi, and Salva, 2014) and chicks (Salvo et al., 2013) as well as humans, suggesting it taps into a feature of cognition that has a broad evolutionary basis.

Participants had to pick up the left-hand disc if the two discs appeared equal in size and the right-hand disc if they appeared different, for many different sizes and apparent sizes of discs. The aperture of their finger–thumb grip was measured as they did so, allowing motor performance and a perceptual decision to be measured in the same task. Participants saw the usual size illusion (as shown by their choice of disc), but their grip fitted the actual disc. Apparently the visuomotor system was not fooled, although the perceptual system was. The same illusion has since been demonstrated in the tactile modality, with blindfolded participants exploring foam cut-outs of the circles using touch (Ziat et al., 2014). And other work with healthy participants has shown that modifying the subjective experience of masked stimuli can be done without changing on the motor effects of those stimuli (Vorberg et al., 2003).

These studies underline the important difference between processing for perception and processing for motor control. The distinction makes sense in evolutionary terms because the constraints on the two systems are different. Fast and accurate responses to changing visual stimuli are essential for catching prey, avoiding dangers, and even basic abilities like standing upright. By contrast, object identification can wait. Rich detail rather than speed may be more important when planning future actions and making strategic decisions, and this

may explain why we have these two different visual systems. The result is that a great deal of what we do is done quickly and accurately, and independently of what we consciously perceive.

Can we now conclude that one of the streams is conscious while the other is a zombie, as Ramachandran and Blakeslee (1998) claim? Although Milner and Goodale were initially cautious about making this distinction, they subsequently took the same view: while the ventral stream provides 'a conscious representation of the world' (Goodale, 2007, p. 626) and 'the sole route to phenomenal visual consciousness' (Milner, 2008, p. 177), 'the visual products of dorsal stream processing are not available to conscious awareness – [. . .] they exist only as evanescent raw materials to provide the unconscious moment-to-moment sensory calibration of our movements' (Milner, 2012). Similarly, Nancy Kanwisher (2001) suggests that 'the neural correlates of the *contents* of visual awareness are represented in the ventral pathway, whereas the neural correlates of more general-purpose *content-independent* processes [. . .] are found primarily in the dorsal pathway' (p. 98). Note that these formulations make several assumptions: that consciousness has contents, that the contents are representations, and that there is a difference between some areas or processes that are conscious and those that are not. We have begun to question every one of these assumptions.

The principles of embodied, or distributed, cognition encourage us to question the very distinction between vision for (conscious) perception and for action. On this view, perception always happens (to a greater or lesser extent) through and for action. Hurley's (1998) book *Consciousness in Action* and Alva Noë's (2005) book *Action in Perception* both understand conscious experience as dependent on, or even constituted by, embodied action, whether actual or potential. More recent experiments with patient D.F. support this view. German psychologist Thomas Schenk found that she can accurately grasp reflections of objects, but only when there is an actual object there to be grasped. That is, her apparently intact vision-for-action actually relies on haptic (touch-based) feedback from objects in the environment to scale her reaching actions and grip size to those objects. This converges with findings from visually normal participants in Schenk's experiments and others (e.g. Bingham, Coats, and Mon-Williams, 2007): for example, if the size of a grasped object is changed while a reflected visual object stays the same size, people are very well able to recalibrate their reach actions and grip sizes using haptic feedback.

These debates, about perception and action, consciousness and unconsciousness, have never been more heated than in discussion of a strange phenomenon known as blindsight.

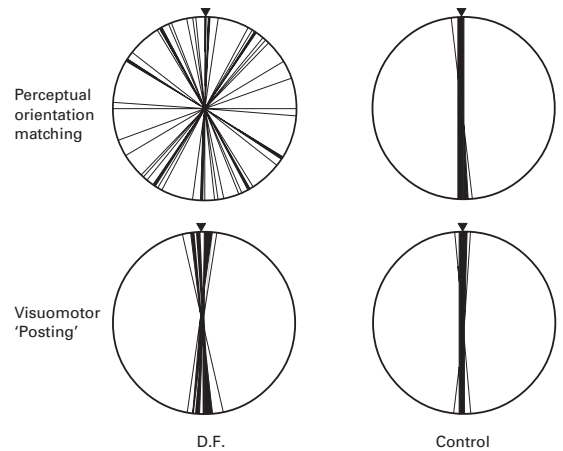


FIGURE 8.10 • Polar plots illustrating the orientation of a hand-held card in two tasks of orientation discrimination, for D.F. and an age-matched control. On the perceptual matching task, both were required to match the orientation of the card with that of a slot placed in different orientations in front of them. On the posting task, they were required to reach out and insert the card into the slot. The correct orientation has been normalised to the vertical (Milner and Goodale, 1995).



FIGURE 8.11 • Diagram showing the 'Titchener circles' illusion. In the top figure, the two central discs are of the same actual size, but appear different; in the bottom figure, the disc surrounded by an annulus of large circles has been made somewhat larger in size so as to appear approximately equal in size to the other central disc (Milner and Goodale, 1995).



FIGURE 8.12 • Which way do the stripes go? When such a display was shown to the blind field of a person with hemianopia (blind on one side), he said he could see nothing at all. Yet when pressed to guess he was able to discriminate vertical from horizontal stripes with over 90% accuracy. This is how the term 'blindsight' originated.

BLINDSIGHT

Imagine the following experiment. A patient, D.B., has had a small non-malignant brain tumour removed from area V1 and this has left him blind on one side. If he looks straight ahead and an object is placed on his blind side, he cannot see it.

In the experiment, D.B. is shown a circle filled with black and white stripes in his normal field. Naturally enough, he says he can easily tell whether the stripes are vertical or horizontal. Now he is shown the same thing in his blind field. He says he can see neither the circle nor the stripes, for he is blind there. Even so, the experimenters encourage him to guess which way the stripes go. He protests that this is pointless, because he cannot see anything, but nevertheless he guesses. He is right 90 or 95% of the time.

'Blindsight' is the oxymoronic term invented for this condition by Oxford neuropsychologist Lawrence Weiskrantz. Together with his colleague Elizabeth Warrington, he tested D.B. from the early 1970s for ten years or more (Weiskrantz, 1986, 1997, 2007). Since then many other blindsight patients have been tested, the most famous of whom is G.Y., who suffered traumatic head injury in a car accident when he was eight years old. Most 'blindseers' have extensive damage to visual striate cortex on one side, which causes degeneration of cells down through the lateral geniculate and even to the retina, while other, non-cortical visual pathways are left intact. Related phenomena such as 'deaf hearing', 'blindsmell', and 'numbsense' have added to the cases in which people deny having conscious sensory experiences and yet behave as though they can see, hear, smell, or feel.

Blindsight seems to be tailor-made for resolving philosophical arguments about consciousness. Yet it has not done so. Blindsight has been used to support qualia and to reject them, to bolster zombies and to undermine them, and to support controversial distinctions between different kinds of consciousness (Dennett, 1991; Block, 1995; Holt, 1999). The arguments have been so long and fierce that it is worth considering blindsight in some detail.

Superficially, the most obvious interpretation goes something like this:

The blindseer has vision without consciousness. He is an automaton or a partial zombie who can 'see' functionally but has none of the visual qualia that go with normal seeing. This proves that consciousness is something separate from the ordinary processes of vision. It proves that qualia exist and functionalism is wrong.

If it were valid, this line of reasoning would have many other implications. For example, it would hold out the hope of finding the place in the brain where 'consciousness happens', the place where visual qualia are produced, or where representations 'enter consciousness'. We would know, for example, that qualia happen in V1 while all the rest of vision goes on elsewhere.

This would encourage speculations about the evolution of consciousness; for if we have qualia *as well as* vision, then there must be some extra function for consciousness.

But this apparently natural way of thinking about blindsight walks straight into all the usual troubles we have met before: the Cartesian theatre where consciousness happens, the Cartesian materialist idea of a 'finishing line' marking entry into consciousness, the hard problem of how subjective qualia can be produced by objective brain processes, and the magic difference between areas or processes that are conscious and those that are not. This explains why blindsight has become such a *cause célèbre*. Either it really has all these dramatic and mysterious consequences and they need explaining, or there is something wrong with the 'obvious' interpretation.

Strong arguments have been made that blindsight does not really exist, that blindsight is only degraded normal vision, and that blindseers are just overly cautious about saying they can see something (Campion, Latto, and Smith, 1983, with peer commentaries; Kentrige and Heywood, 1999). All these arguments have been effectively countered, but there is a grain of truth to them. Although standard blindsight is a severely impoverished form of sight, blindseers are sometimes aware of certain kinds of stimuli in their blind field, especially fast-moving, high-contrast ones. This residual ability makes sense in anatomical terms because there is a minor visual pathway that bypasses V1 and has projections to V5, which is motion-sensitive. Indeed, activity in V5 has been shown in G.Y. by PET scan (Barbur et al., 1993). In one experiment (Morland, 1999), G.Y. was asked to match the speed of moving stimuli shown in his blind field to those in his seeing field. The results showed that as far as motion is concerned, his perception is the same in both. Yet he did not identify the experience as really 'seeing', and explained that it was difficult to know *how* to describe his experience: 'the difficulty is the same that one would have in trying to tell a blind man what it is like to see' (Weiskrantz, 1997, p. 66). This makes sense: it is very difficult to imagine what it is like to see movement without seeing the thing that is moving, yet that is the ability G.Y. has. British psychologist Tony Morland concludes that primary visual cortex is not needed for consciousness, but it is needed for binding the features of objects. So the experience of movement in blindsight is just that – seeing movement that is not bound to a moving object.

Some blindseers also use appropriate eye movements to track moving objects they cannot see, or mimic the path of an invisible stimulus with their hands. Some can make reasonably accurate movements to grasp invisible objects, and even to post invisible cards through slots with the correct orientation. Odd as this seems, it makes sense in terms of the distinction between the



FIGURE 8.13 • The person with blindsight has to be pressed to guess the orientation of a line he cannot see. Yet his guesses can be very accurate. Is he a partial zombie who has vision without conscious vision?

Does blindsight provide 'a case where all the functions of vision are still present, but all the good juice of consciousness has drained out? It provides no such thing.'

(Dennett, 1991, p. 325)

‘the difficulty is the same that one would have in trying to tell a blind man what it is like to see’

(G.Y.)

dorsal and ventral streams. Milner and Goodale suggest that ‘blindsight is a set of visual capacities mediated by the dorsal stream and associated sub-cortical structures’ (1995, p. 85). This fits with Weiskrantz’s observation that ‘the intact field seems to be biased towards object identification, and the blindsight field towards stimulus detection’ (1997, p. 40). If this is correct, it means that the detection of stimuli in blindsight is based on visuomotor responses.

Milner and Goodale also note that G.Y. reports a different non-visual experience when asked to use different visuomotor responses. They do not conclude that consciousness is obliterated along with the ventral stream, but rather that there may be ‘a distinct non-visual experiential state associated with each different visuomotor system activated’ (1995, p. 79). In their view, blindsight should be understood not as perception without consciousness, but – like visual form agnosia – as action without perception.

Further evidence comes from studies of sensory substitution, in which people are given information in one sense to replace another, for example touch or sound to replace vision (Concept 8.1). They too have trouble describing what the experience is like, but with practice it comes to seem more and more like seeing. If this is correct, it suggests that perceptual consciousness is part of learning a new sensorimotor skill, rather than being something separate from it.

This suggests that the arguments about blindsight may all start from a mistaken premise. Maybe they treat as mysterious, paradoxical, and implausible something that in fact is not. For Daniel Dennett, there is no categorical difference between what blindseers do with visual information and what the rest of us do. To explain why not, Dennett notes that in most experiments blindseers have to be prompted to guess, and are given no immediate feedback on their success. Dennett now imagines training a blindsight patient by giving him feedback on his guesses, until he comes to realise that he has a useful ability. Next he is trained, again by giving feedback, to guess on his own, without being prompted. After this training he should spontaneously be able to talk about, act upon, and use the information from his blind field just as well as from his seeing field. Others have dubbed this ‘super blindsight’ (Block, 1995; Holt, 1999) and it has been much disputed.

The argument really hinges on this question – if the super-blindseer could really use the information



CONCEPT 8.1

SENSORY SUBSTITUTION

Can a person who is blind learn to see? Retinal implants are available and completely artificial eyes may one day be possible, but for now the task of wiring them into the brain is too difficult. Another way of solving the problem is to substitute one sense for another.

The first attempts at sensory substitution were made by Paul Bach-y-Rita in the late 1960s (Bach-y-Rita, 1995). Signals from low-resolution cameras on special glasses went to an array of just 16 by 16 vibrators on the blind person’s back. Even with this crude device people could walk about, read signs, and even identify faces. Much higher-resolution devices followed (called Tactile Vision Substitution Systems, TVSS), with tactile arrays on the back, abdomen, thigh, and fingertips. After sufficient training with TVSS, blind people experienced the images as being out in space rather than on their skin, and learned to use parallax, depth, looming, and other visual cues.

about a stimulus in his blind field in this way, would that mean he was necessarily conscious of it? It is worth giving this thought experiment careful consideration to decide on your own answer. Functionalists would say yes (because being conscious *is* performing these functions) and proponents of enactive and sensorimotor cognition probably would, too (because consciousness is constituted by our interactions with the physical world), while those who believe in the existence of qualia, conscious inessentialism, and the possibility of zombies would say no (because the functions and the qualia are separate things and the super-blindseer has one but not the other).

Maybe by taking this route the mystery of blind-sight starts to disappear. Imagine that as the super-blindseer became better and better trained, he would stop denying having qualia, because his experiences would match the quality of the abilities he came to have. If he could be trained to act on and talk about – in Block’s terms, to have access to – stimuli in his blind field, then he would, by definition, also become conscious of them. Interestingly, there is now evidence that through neural plasticity and practice, people with cortical blindness gradually regain some conscious vision in the blind field (Melnick et al., 2016).

Weiskrantz suggests that blindseers lack what he calls the ‘commentary stage’ in which information becomes available for comment, either verbally or in other ways. So, again, the super-blindseer who could comment on his own abilities would thereby become conscious of them. This is similar to HOT theory, in which information is conscious only if there is a higher-order thought to the effect that the person is experiencing it.

But this ‘neural monitor’ is, according to some, ‘no more than a fanciful expedient designed to explain away the paradox of blind-sight’ (Bennett and Hacker, 2003, p. 396). Indeed, there is really nothing paradoxical about the phenomena of blindsight; the paradox is created by the confused ways in which neuroscientists try to describe them (using terms like blindsight or unconscious awareness). Part of the trouble is that we want to ask whether the blindseer *really sees* or not, and

Because the tongue is far more sensitive than the back, other interfaces involve gold-plated electrodes on the tongue. By moving the video camera around, the user can explore the environment as sighted people do by moving their eyes. The effects are dramatic. One blind man even climbed Everest using this technology. Within a few hours, one congenitally blind woman was able to move around, grasp objects, and even catch and toss a ball. She specially asked to see a flickering candle – something she had never been able to experience through any other sense (Bach-y-Rita and González, 2002).

A similar array on the tongue was used to replace vestibular feedback in a woman who had lost her vestibular system and could not even stand upright on her own. Using the new system she could stand almost immediately, without any training.

In a completely different approach, sound is used to replace vision. In Peter Meijer’s (2002) method, a video image is converted into ‘soundscapes’: swooping noises that act like sound-saccades, in which pitch and time are used to code for left–right and up–down in the image. Meijer put the necessary software on the web, and among those who tried it was Pat Fletcher (2002), who was blinded in an industrial accident in 1999. The system took her many months to master, unlike the tactile systems, but eventually she began to see depth and detail in the world.

But is it really vision? Fletcher says it is, and she does not confuse the soundscapes with other sounds. She can have a conversation with someone while using the soundscapes to look at them, and she even dreams in soundscapes. But it is not clear how ‘visual’ these experiences really are, and some have likened sensory substitution to an acquired synaesthesia (Ward and Wright, 2014).

All this has profound implications for the nature of sensory awareness. The ease with which one sense can stand in for another suggests that there is nothing intrinsically visual about information that comes through the eyes, or intrinsically auditory about information coming through the ears. Rather, the way the information changes with a person’s actions is what determines how it is experienced. This fits well with sensorimotor theory, which treats vision and hearing as different ways of interacting with the world. The same conclusion is reached from experiments

in which the sensory systems of ferrets are rewired soon after birth. If visual information is routed to auditory cortex, that cortex develops orientation-selective responses, maps of visual space, and control of visual behaviour as visual cortex normally would (Sur and Leamey, 2001). In other words, it seems as though the nature of the input helps structure sensory cortex.

These kinds of research might help solve a classic mystery: how the firing of some neurons leads to visual experiences while identical kinds of firing in different neurons leads to auditory experiences (O'Regan, 2011). Perhaps more important for people who are blind, it suggests that seeing does not necessarily need eyes.



FIGURE 8.14 • Pat Fletcher, shown here with Peter Meijer (left) and David Chalmers (right), is seeing with 'Soundscapes', also known as 'The vOICe' (get it?). She wears headphones and has tiny video cameras concealed in her glasses. A notebook computer in her rucksack carries out the video-to-audio transformations that enable her to see well enough to walk about, pick up objects, and even recognise people. But is it seeing? She says it is.

'Blindsight is paradoxical only if one regards vision as a unitary process'

(Milner and Goodale, 1995, p. 86)

that question cannot be answered because seeing is not all or nothing but depends on the normal convergence of affect and bodily behaviours (including verbal responses), and in blindsight this normal convergence is disrupted. Hence we end up not knowing whether *see* is the right word to use. But maybe our words and our definitions are the problem.

Milner and Goodale agree. 'Blindsight is paradoxical only if one regards vision as a unitary process' (Milner and Goodale, 1995, p. 86). In fact, there is no single visual representation that is used for all purposes, but lots of semi-independent subsystems like those in the ventral and dorsal streams. The mystery of consciousness does not disappear but looks quite different for those who abandon the idea of unified consciousness, a single picture in the mind, a show in the Cartesian theatre, or 'a bogus concept of introspection and privileged access' (Bennett and Hacker, 2003, p. 396).

Maybe we need to remember that for all their differences in emphasis, the two streams are both part of a single system: 'perception' and 'action' are not neatly separable, and when perceptual information of one kind is lacking, it can be supplemented by another kind (Wilson, 2012). This ties in with arguments about the sensorimotor basis of vision (Chapter 3) and with what we know about neural plasticity and sensory substitution. Clean distinctions between brain areas responsible for x, y, and z are always tempting, and are encouraged by research involving individuals with unusual kinds of brain damage affecting specific areas, as well as by the kinds of technology we use to investigate the brain (Chapter 4). But if we jump to the conclusion that some areas are responsible for conscious experience and others for unconscious processes, the water gets murkier still.

Before we turn to one last example of the claimed power of consciousness, we can return to answer our questions about that scrumpled paper ball. The findings we have surveyed suggest that conscious perception of the ball depends on processing that is separate from and too slow to play a role in guiding the fast catch. So, although the causal sequence *seems* to be 1) consciously perceive, and 2) act on the basis of conscious experience, we now know that it cannot be.



PRACTICE 8.2

WAS THIS DECISION CONSCIOUS?

Going about your ordinary activities, you make countless large and small decisions, from exactly where to put your foot as you walk upstairs, to where to go for your holiday, or whether to take that job. But perhaps it might be more accurate to say that your whole body is making decisions, rather than that 'you' are. Watch these decisions as they happen, and for each one that you notice ask yourself '**Was this decision conscious?**' As you begin to notice more and more decisions being made, what happens? Is it obvious which are made consciously and which unconsciously? Are there certain types of decision that are more often conscious? Does anything happen to your sense of agency? What?




FIGURE 8.15 • Super-blindsight. Imagine that a person with blindsight is trained to make spontaneous guesses about things he cannot see.

INTUITION AND CREATIVITY

Intuition is often thought to be strange, inexplicable, or even paranormal, but need it be?

There are at least three components to intuition. First are the cognitive processes in which the brain extracts information from complex patterns to guide behaviour, such as when finding your way around new software or guessing which queue will be shortest in the supermarket. Second are all the social skills and implicit impressions which we cannot articulate or formalise, from the 'feeling' that someone is untrustworthy, to judging the best time to break bad news to a friend. These have tended to be undervalued in comparison with explicit, intellectual skills, perhaps because children readily pick them up and adults do not appreciate the complexity of what is involved. Take the example of judging someone untrustworthy. This may depend on long years of meeting people who look, stand, move their eyes, and twitch other muscles in different ways, and then noting (quite unconsciously) whether they kept their word or not. None of us can explain how we do this, whether we find it easy or a real struggle. People with autism, for example, may find it difficult to understand complex social emotions or adopt other people's points of view, while those with social anxiety disorder may compulsively overinterpret their own and others' words and actions, usually in negative ways.

The notion of 'women's intuition' is sometimes laughed at, but women may be more intuitive, in this sense, because they generally have better verbal skills, are more interested in relationships, and gossip more about social matters than men

'The mechanisms of consciousness are also embodied in our comportment within the (social) world, and not just limited within our brain'

(Froese et al., 2014, p. 8)

'emotions and feelings may not be intruders in the bastion of reason at all: they may be enmeshed in its networks, for worse and for better'

(Damasio, 1994, p. xxii)

'all the contents of consciousness are in harmony with each other, and with the goals that define the person's self'

(Csikszentmihalyi and Csikszentmihalyi, 1988, p. 24)

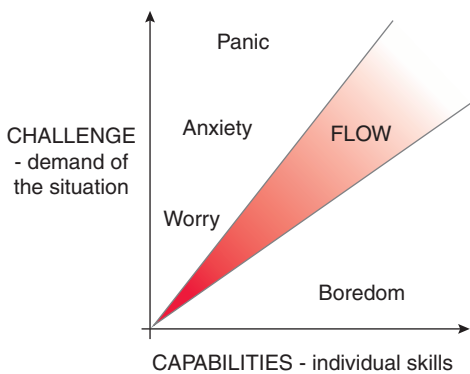


FIGURE 8.16 • According to Csikszentmihalyi, the state of flow occurs when the challenges presented by a task are proportional to the person's capabilities, thus avoiding both boredom and anxiety.

do. So, if they spend a lot of time soaking up the covariations in our vastly complex social world, they may more often be right when they say 'I don't trust that man' or 'I think those two are falling in love', even if they cannot articulate the reasons for their judgement.

The third component, though not separate from the others, is emotion, as when people say 'it just felt wrong' or 'I just knew it was the house for me'. Although emotion and reason have traditionally been opposed, they are equally integral to a process which helps flexibly guide appropriate actions (Frijda, 2007). Portuguese neurologist Antonio Damasio (1994) is famous for arguing that reason cannot operate without emotion. He studied many patients with frontal lobe damage who became emotionally flat, yet far from turning into super-rational decision-makers, they became paralysed with indecision, every little choice becoming a nerve-wracking dilemma. They could still rationally compare alternatives but lacked the feelings that make decisions 'seem right'. This implies that *Star Trek's* Spock would not be the impressive Starfleet first officer he is portrayed as, for suppressing his feelings in favour of logic would make him unable to decide whether to get up in the morning, when to speak to Captain Kirk, or whether the Klingons are bluffing.

This interpretation needs caution, though, because the fact that frontal lobe damage affects both emotion and decision-making does not prove that emotion is needed for decision-making; both might depend on some other affected capacity, for example.

Creativity might also entail these explicit and intuitive skills coming together to generate new insight. Many creative writers, thinkers, scientists, and artists claim that their best work just 'comes' to them. They have no idea how they do it, and may feel as though the poem, painting, or solution to the scientific problem just shaped itself without their conscious effort or awareness. Creative people tend to score high on measures of imagery, fantasy-proneness, hypnotisability, and 'absorption'; that is, they can easily become so absorbed in a book, film, or their work that they are oblivious to everything else. Some describe this timeless feeling of total immersion as a selfless state of 'flow' (Csikszentmihalyi, 1975; see our website for more detail). Finding flow

depends on getting the right balance between the challenge you face and the skills you bring to tackling it. When a challenge is too great, anxiety results; when too slight, boredom sets in. But when challenges and skills are perfectly matched, flow can take over. Although flow is usually described as a state of consciousness, it might better be described as a state in which the distinctions between conscious and unconscious processing disappear. All of a person's skills are called upon, and there is no longer any self to say just what 'I' am conscious of.

Creativity also often involves working hard on a problem and failing to solve it. Then, after resting or doing something else, the solution just 'pops into mind'. The hard work is essential but so are the unconscious processes, and these need time and leaving alone. This process, called incubation, complicates the simple fast/slow distinctions just discussed, since it seems to rely on conscious effort, extended unconscious processing, and then a sudden moment of inspiration. Studying incubation in the real world is difficult, but tricky puzzles and devious brain-teasers may provide something of the same effect.

There are many claims of a cosmic creative force, or a power of consciousness beyond the mind that is deeply mysterious and beyond the reach of science, but there is a well-known cosmic force that really is creative: the evolutionary algorithm (Chapters 10 and 11). Could cultural evolution be a force for human creativity, with ideas emerging from the process of copying previous ideas, but with variation and selection?

'it is sometimes a good idea to pull off the Information Super-Highway into the Information Super Lay-By'

(Claxton, 1997, p. 14)

Thinking about creativity in this way means seeing individual creators in their social and intellectual context. When James Watt was worrying about heat loss, it was because he had seen steam engines, and knew about the manufacturing processes of his time. When inspiration came to Szilard, he was deeply immersed in the atomic science of his day; when Coleridge fell into his sleep he had just been reading a book about the palace built by the Khan Kubla. In other words, they had been soaking up the memes of the culture around them.

On this meme-based view, what makes creative people unique is how they recombine old memes to make new ones and have the intuition to feel which of the billions of possible combinations is worth pursuing. The individual person is an indispensable part of the creative act, but the real driving force is cultural evolution.

Don't even try to talk about the learning curve. Don't bother citing the months of deliberate practice that precede the unconscious performance, or the years of study and experiment leading up to the gift-wrapped Eureka moment. So what if your lessons are all learned consciously? Do you think that proves there's no other way? Heuristic software's been learning from experience for over a hundred years. Machines master chess, cars learn to drive themselves, statistical programs face problems and design the experiments to solve them and you think that the only path to learning leads through sentience? You're Stone-age nomads, eking out some marginal existence on the veldt – denying even the possibility of agriculture, because hunting and gathering was good enough for your parents.

(Peter Watts, *Blindsight*, 2006)

The other crucial context is the physical world. Some philosophers, most famously Andy Clark, believe that from our notebooks to our smartphones and GPS, the objects we use are part of our cognitive architecture. Linking cultural evolution with the extended mind is the idea of the cognitive niche. In biology, niche construction occurs whenever species act on their environments in ways that change the factors that will be adaptive in the future (Wheeler and Clark, 2008). By building a web, a spider changes the sources of natural selection within its niche; a beaver's evolved niche includes the dam constructed by its parents and the changes in the flow of the river caused by the dam. Niche-construction leads to new feedback cycles for humans, too: think of the novice bartender who inherits a highly structured environment to help her learn how to remember drinks orders and serve them swiftly.

'unravelling the complex interplay between genes, environments and embodied action [...] will surely be one of the great intellectual adventures of the 21st century'

(Wheeler and Clark, 2008, p. 3572)

Human niche construction shows the physical and cognitive worlds to be inseparable. Language is probably the most powerful example of an ecology that shapes human development, at both individual and species level – it shapes our ways of thinking and acting, and everything that we think of as our human nature, including consciousness itself. Maybe language makes thinking and acting feel different enough that we have learned to call the kinds of thinking and acting we know how to talk about *conscious*, when in fact speaking and writing are just kinds of action, not fundamentally different from pressing a button, catching a ball, moving our eyes, or opening and closing our fingers under general anaesthetic.

THE UPSHOT

The evidence for unconscious perception and action, and for intuitive decision-making and creativity, shows that some popular ideas about consciousness have to be wrong. To make this clear, let's consider three basic ways of thinking about consciousness.



PROFILE 8.2

Andy Clark (b. 1957)



Andy Clark wants to extend our notion of mind far beyond our brains and even our bodies. Human minds can be 'extended minds', he says, realised by neuronal, bodily, and even technological elements, such as smartphones and good old-fashion-

ed pencil and paper. He believes that the drive towards cognitive extension is so deeply ingrained that we are natural-born cyborgs: beings whose minds and selves arise at the changing intersections between biology and technology. More recently he has come to believe that work on the 'predictive brain' – with feedback between top-down expectations and bottom-up inputs – holds the key to the delicate dance between brain, body, and world. He loves electronic music (especially old-school techno), American comics, and pulp detective fiction. He owns a 47 ft houseboat, *Love and Rockets*, named after the comic and decorated (like his own body) with tattoo art. He has held posts in St Louis and Bloomington, and now holds the Chair in Logic and Metaphysics at the University of Edinburgh.

The first is the traditional (ever-tempting) idea of a Cartesian theatre. Consciousness is like a multisensory cinema with information coming into consciousness, for 'me' to experience and act upon. In its most extreme view, this assumes that sensory information can lead to action only once it has 'become conscious' in the Cartesian theatre. We had already found many reasons for rejecting this view, and the phenomena explored in this chapter provide more.

The second view allows for unconscious perception and learning, but still fails to throw out the theatre metaphor. The idea is something like this: sensory information enters the system, whereupon two distinct things can happen to it. Either it goes into consciousness and is acted upon consciously, or it bypasses consciousness and is acted on unconsciously, perhaps by using routes through the brain that lead to motor output without ever actually 'reaching consciousness'.

This second theory, a form of Cartesian materialism, is probably the most common in consciousness studies today. While rejecting the notion of a homunculus watching events on a mental screen, it retains the essential idea that things are either 'in' or 'out' of consciousness. As we have seen, phrases such as 'enters consciousness', 'available to consciousness', or 'reaching consciousness' tend to imply such a theory. The tricky issues surrounding the border between 'in' and 'out' of consciousness are sometimes dealt with by proposing a 'fringe' consciousness (e.g. Baars, 1988), or by avoiding 'fuzzy' cases such as blindsight or subliminal perception.

The findings discussed above suggest a more radical third theory. To recap: thresholds of conscious experience are not

fixed but depend on variable response criteria; there is no undisputed measure for deciding whether or not a stimulus has been consciously perceived, or an action consciously carried out, or a skill consciously learned, or a decision consciously made; there are many examples where the answer – conscious or unconscious – depends on the way you ask the question. All this threatens the idea which seems so intuitively obvious: that a given stimulus is unequivocally either ‘in’ or ‘out’ of consciousness, or that a given physical or cognitive act is unequivocally performed consciously or not. Indeed, it is not clear that there is any coherent way to argue for a theory of conscious causation.

Instead, this third way suggests that sensory information is processed in multiple ways, with different consequences for different behaviours, and that the type of behaviour involved may affect how the processing happens from the very beginning (Marcel, 1983). Some of these behaviours are usually taken as indications of consciousness, such as verbal reports or choices between clearly perceptible stimuli, while others are usually considered to be unconscious, such as fast reflexes, or guesses. In between lie many behaviours which are sometimes taken to indicate consciousness and sometimes not. *But there is no right answer.*

We have considered a wide range of evidence in this chapter, from the ability to distinguish between near-identical weights by judging with one’s finger the pressure they create on a weighing scale (Peirce and Jastrow, 1885) to the capacity to hold a conversation using one’s forearm while under anaesthetic (Alkire, Hudetz, and Tononi, 2008), to the lightning-quick appraisals by which, in every social encounter, we distinguish friend from foe. The range of experimental measures includes every imaginable combination of button-pressing, rating scales, and verbal reports – so many that maybe the field needs to develop better ways of comparing results (Rothkirch and Hesselmann, 2017). What do any of these correlations tell us about whether or not the perception, the action, or the person was ‘really conscious’? Who gets to decide that a movement of the finger means unconscious, whereas movements of the lips mean conscious (unless you insist you’re just guessing)?

On this third way of thinking, nothing is ever ‘in’ or ‘out’ of consciousness, and phrases such as ‘reaching consciousness’ or ‘available to consciousness’ are either meaningless or are a short-hand for ‘leading to verbal report’ or ‘available to influence behaviours taken to indicate consciousness’. Calling an action or a perception conscious is another example of the mereological fallacy. An event itself is never either conscious or unconscious, but they can be more or less likely to lead the *person* to say they were conscious of it.

The major problem here is what people say about their own experience. Many people say that they know for sure what is in their consciousness and what is not, even if they cannot always explain what they mean. One way out is to take the intuitions seriously, but to accept that they are illusions and then try to explain how the illusions comes about (Dennett, 1991; Blackmore, 2016; Frankish, 2016b). But the gulf between the evidence and the intuition is a familiar one, and just one more reason why the problem of consciousness is so perplexing.

WAS THIS DECISION CONSCIOUS?

‘the limits on unconscious processing are set by the means by which the stimuli are rendered consciously inaccessible’

(Kihlstrom, 1996, p. 39)



FIGURE 8.17 • Solution to 8.14

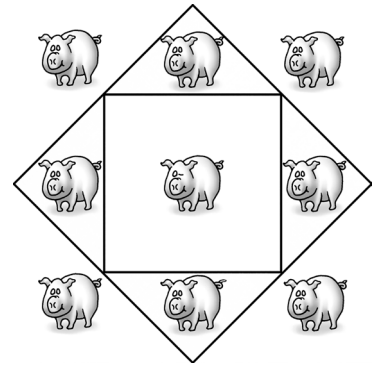


FIGURE 8.18 • Solution to 8.15, and did you notice the wink?

READING




FIGURE 8.19 • Solution to 8.16

Dijksterhuis, A., Aarts, H. and Smith, P. K. (2005). The power of the subliminal: On subliminal persuasion and other potential applications. In R. R. Hassin, J. S. Uleman, and J. A. Bargh (Eds), *The new unconscious* (pp. 77–106). Oxford: Oxford University Press.

Sets out the evidence for, and the theoretical implications of, the limitations of 'conscious thought'.

Frankish, K., and Evans, J. St B. T. (2009). The duality of mind: An historical perspective. In K. Frankish and J. Evans, *In two minds: Dual processes and beyond* (pp. 1–29). Oxford: Oxford University Press.

A history of how the mind has been divided in two, culminating in modern dual-process theories.

Kentridge, R. W. and Heywood, C. A. (1999). The status of blindsight. *Journal of Consciousness Studies*, 6, 3–11.

A summary of debates about what blindsight is and means.

Kihlstrom, J. F. (1996). Perception without awareness of what is perceived, learning without awareness of what is learned. In M. Velmans (Ed.), *The science of consciousness* (pp. 23–46). London, Routledge.


Traces the history of research on dissociations of consciousness from perception and learning.

Merikle, P. (2007). Preconscious processing. In M. Velmans and S. Schneider (Eds), *The Blackwell companion to consciousness* (pp. 512–524). Oxford: Blackwell.

Uses vision as a case study for distinguishing conscious from unconscious cognitive processing, with particular focus on questions about measurement.

Velmans, M. (2002). How could conscious experiences affect brains? *Journal of Consciousness Studies*, 9(11), 3–95.

Proposes a new account of mental causation to address this question, trying to avoid the problems of reductionist (mind=brain) and non-reductionist accounts with a dual-aspect theory. Velmans's response to the commentaries introduces reflexive monism (Chapter 17).



Agency and free will

CHAPTER

NINE

'We know what it is to get out of bed on a freezing morning in a room without a fire, and how the very vital principle within us protests against the ordeal,' said William James, describing the agonising, the self-recrimination, and the lure of comfort against the cold. 'Now how do we *ever* get up under such circumstances?' he asked. 'If I may generalize from my own experience, we more often than not get up without any struggle or decision at all. We suddenly find that we *have* got up' (James, 1890, ii, p. 524). When the inhibitory thoughts briefly cease, he said, the idea of getting up produces its appropriate motor effects, by 'ideo-motor action', and we are up. What, then, is the role of free will?



PRACTICE 9.1

AM I DOING THIS?

When you find yourself asking 'Am I conscious now?' observe what you are doing and ask yourself '**Am I doing this?**' You might be walking, drinking a cup of coffee, or picking up your phone to ring a friend. Whatever it is, ask yourself what caused the action. Did you consciously think about it first? Did your own conscious thoughts cause it to happen? Did it just happen by itself?

You might like to take a short time – say ten minutes – and try to observe the origins of all your actions during that time. In each case ask, 'Did **I** do that?'

The problem of free will may be the most-discussed philosophical problem of all time, going back at least to the Greek philosophers two thousand years ago (for an overview of psychology and free will see Baer, Kaufman, and Baumeister, 2008). The basic question is whether or not we are free to choose our actions and make decisions. The question for us here is whether consciousness has any role to play in our acting freely, or feeling that we do.

Major religions, especially the monotheistic religions, depend heavily on the belief that we do have freedom of will. Christianity teaches the doctrine of original sin and that God gave us the choice between good and evil, between His way and that of the Devil. Islam teaches that we are accountable for every choice, despite the fact that Allah already knows everything that will happen. If asked *who* has this choice, or *who* is good or evil, believers will point to the human soul or spirit – that immaterial, conscious being which is ultimately responsible and after death is rewarded or punished in heaven or hell. Arguably these threats and promises are meme-tricks which benefit religious memes by keeping people hopeful or afraid (Dawkins, 1976; Blackmore, 1999). Certainly the belief in free will is intrinsic to these religions, and it is also widespread across many cultures (Sarkissian et al., 2010).

There are two main problems: one for free will, the other for its absence. The first is determinism: if this universe runs by deterministic laws, then everything that happens must be inevitable, so the argument goes, and if everything is inevitable, there is no room for free will; no point in my 'doing' anything; no sense in which I 'could have done otherwise'. The second is moral responsibility: if I am not truly free to choose my actions, then how can I be held morally or legally responsible for them?

This is the excellent foppery of the world, that when we are sick in fortune (often the surfeits of our own behaviour) we make guilty of our disasters the sun, the moon, and stars: as if we were villains on necessity; fools by heavenly compulsion; knaves, thieves, and treacherous by spherical predominance; drunkards, liars, and adulterers by an enforced obedience of planetary influence; and all that we are evil in, by a divine thrusting on. An admirable evasion of whoremaster man, to lay his goatish disposition on the charge of a star!

(Shakespeare, *King Lear*, I.ii, 1606)

Determinism may or may not be true, but note that it is not the same as predictability. For instance, physicists describe random events such as radioactive decay which seem both undetermined and unpredictable, and chaotic processes which are determined but unpredictable, for example because of sensitive dependence on starting conditions. Whether the universe is really deterministic is a different argument from whether determinism is compatible with free will. Among modern philosophers, non-compatibilists argue that if the universe is deterministic, then free will must be an illusion, while compatibilists find many and varied ways in which determinism can be true and yet free will remain free, for example by stressing ways in which it can still be true that 'I could have done otherwise'.

• SECTION THREE: BODY AND WORLD

There are many arguments here, and little agreement, except perhaps for a widespread rejection of free will as a magical or supernatural force. If free will exists, it is certainly not magic. The question is, what other possibilities are there? If we add chance or randomness, as modern physics does, we get back to the Greek philosopher Democritus, who is reputed to have said that 'everything in the universe is the fruit of chance and necessity'. And it is not chance or randomness that believers in free will want, but some way in which their own efforts really make a difference.

[S]cience itself will teach man [. . .] that he does not have and, in fact, has never really had any caprice or will of his own, and that he himself is something like a piano key or the stop of an organ, and that there are, besides, things called the laws of nature; so that everything he does is not done by his willing it, but is done of itself, by the laws of nature. Consequently we have only to discover these laws of nature, and man will no longer have to answer for his actions and life will become exceedingly easy for him. All human actions will then, of course, be tabulated according to these laws, mathematically, like tables of logarithms up to 108,000, and entered in a calendar; or, better still, there would be published certain edifying works, comparable to today's encyclopaedic lexicons, in which everything will be so clearly calculated and explained that there will be no more deeds or adventures in the world.

(Fyodor Dostoyevsky, *Notes from the Underground* [Записки из подполья], 1864, our translation)

This is where the connections with self and consciousness come in, for we feel as though 'I am the one who acts; 'I am the one who has free will; 'I am the one who consciously decided to spring out of bed early this morning. When the chosen action then happens, it *seems* as though my conscious thought was responsible. Indeed, it seems that without the conscious thought I would not have done what I did, and that *I* consciously caused the action by deciding to do it. The question is: does consciousness really play a role in decision-making and choice? Is this sense of conscious agency justified or illusory? We began to tackle this question in Chapter 8 when trying to distinguish conscious from unconscious action, and considered theories that do, and do not, give a causal role to consciousness. Here we will explore how consciousness relates more generally to our sense of personal agency and free will.

As ever, William James got to the heart of the matter when he said:

the whole feeling of reality, the whole sting and excitement of our voluntary life, depends on our sense that in it things are *really being decided* from one moment to another, and that it is not the dull rattling off of a chain that was forged innumerable ages ago. This appearance, which makes life and history tingle with such a tragic zest, *may* not be an illusion.

(James, 1890, i, p. 453)

As we will see (Chapter 16), James rejected the idea of a persisting self but still believed that the sense of effort in both attention and volition was not an illusion but the truly causal force of conscious, personal will (Chapter 7). His interesting use of the passive construction *things are really being decided* (with no one necessarily doing the deciding) perhaps hints at his lifelong ambivalence. By the time he wrote his late work *The Varieties of Religious Experience: A Study in Human Nature*, his view had changed to one in which renouncing all desire and choice is what leads to freedom.

He became, not a man with a mind, but a great instinct. His hands were like creatures, living; his limbs, his body, were all life and consciousness, subject to no will of his, but living in themselves. Just as he was, so it seemed the vigorous, wintry stars were strong also with life. He and they struck with the same pulse of fire, and the same joy of strength which held the bracken-frond stiff near his eyes held his own body firm.

(D. H. Lawrence, *Sons and Lovers*, 1913)

An illusion, we should remember, is not something that does not exist, but something that is not what it seems. So how does it *seem* to you? Does it *seem* as though you have free will, that your decisions are freely made by your conscious mind, even some of the time? If so, then ask yourself whether this could be an illusion, and if it is an illusion, you will need to work out how you can possibly live with that idea (Chapter 18). If it does *not* seem to you that your actions are initiated by conscious decisions, then you may read all this with an air of amused detachment.

Note that we are concerned here with *consciousness*. The question is not whether human beings are agents or can make choices. We may safely assume that they are and can. Humans are living biological creatures that survive, like all other creatures, by having boundaries between themselves and the outside world, and by taking control over certain aspects of that world. They respond to events, make intricate plans with many available options, and act accordingly, at least when not restrained or coerced.

Neither need we doubt that thought, deliberation, and emotions play a part in decisions. Weighing up possible actions and comparing their likely outcomes is what intelligent animals are good at, from a cat deciding when to pounce, to a chimpanzee predicting the likely consequences of challenging a dominant ally.




FIGURE 9.1 • Brain activity preceding a voluntary action of the right hand. The frontopolar cortex (shown in green) forms and deliberates long-range plans and intentions. The pre-supplementary motor area (shown in red) begins the preparation of the action; together with other premotor areas, it generates the readiness potentials (red trace) that can be recorded from the scalp. Immediately before the action takes place, M1 (shown in blue) becomes active. In later stages of preparation the contralateral hemisphere is more active than the ipsilateral hemisphere; this is reflected in a lateralised difference between the readiness potentials that are recorded over the two hemispheres of the brain (solid and dotted blue traces). Finally, neural signals leave M1 for the spinal cord and the contralateral hand muscles. The contraction of the muscles is measured as an electrical signal, the electromyogram (Haggard, 2008).

'Even robots believe they have free will, even if they don't'

(O'Regan, in Blackmore, 2005, p. 172)



ACTIVITY 9.1

Getting out of bed on a cold morning

Try William James's famous meditation (as he called it) and watch what happens when you get out of bed on a cold morning. If you don't live somewhere cold enough, just choose a morning when you really *don't* want to get up. Alternatively try getting out of a bath when the water is going cold and you've been in there too long, or out of a hot shower when you're really enjoying it.

Watch what happens. What thoughts go through your mind as you struggle to get out? What emotions do you feel? Do you speak to yourself or try to persuade yourself? If so, who or what is struggling against whom or what? What happens in the end? You might like to write a short description as James did (1890, ii, p. 524–525).

Comparing descriptions can make for a lively class discussion. What are the implications for free will?

'All theory is against the freedom of the will; all experience is for it'

(Samuel Johnson, 15 April 1778, in Boswell, 1791/1952, p. 393)

We can look to see which parts of the brain and the rest of the body are involved in such decision-making and, in principle at least, trace how they lead to particular decisions and actions. But is this any different from exploring Google's search algorithms to see how it chose which list of links to show me when I asked it 'what is consciousness?'. Google's choices, we may assume, are fully determined by its fiendishly complicated algorithms so it could not, *in those circumstances*, have done otherwise. Does Google need consciousness?

THE NEUROANATOMY OF VOLITION

Consciousness may feel like the cause of voluntary action, but when we look inside the brain, we see lots of areas involved in carrying out the different phases of a voluntary action (see Haggard [2008, pp. 937–938] for simple models of the circuits and the phases). An obvious question to ask is where, if anywhere, does consciousness come in?

An extensive network of brain regions in the medial and lateral frontal cortex, as well as the parietal cortex, are thought to be related to 'internally guided' behaviour (Brass et al., 2013). Multiple areas of neural activity converge on primary motor cortex, which carries out motor commands by sending signals through the spinal cord to your muscles (Spence and

Frith, 1999). There are different pathways for 'internally' and 'externally' triggered actions, although really there is a continuum between the two (Haggard, 2008). Externally triggered actions show activation in the cerebellum and premotor cortex. Intentional actions correlate with activity in prefrontal regions. These include the supplementary motor area (SMA) involved in the sequencing and programming of motor acts to fit a 'motor plan'; the preSMA, which may be the source of the early part of the readiness potential (the activity leading up to the muscle movement); and the anterior cingulate, a complex area involved in emotion and pain as well as attention to, and selection of, information needed for action. There is also Broca's area (in the left inferior frontal gyrus in most right-handed people), which produces the motor output for speech.

Some of this is known from the effects of brain damage – for example, from the famous case of railroad worker Phineas Gage. In 1848 a tamping iron was blown straight through his frontal cortex, leaving him a changed personality and no longer able to behave responsibly (Damasio, 1994). Damage to dorsolateral prefrontal cortex can lead to a lack of spontaneous activity, and to repetitive, stereotypic actions. People with lesions in the preSMA are prone to automatic actions in response to environmental triggers, as though unable to stop themselves eating an apple they see in front of them, or putting on a garment because it's there. Lesions of the prefrontal region and corpus callosum can produce the extraordinary complaint of 'alien hand', in which patients say that their hand has a will of its own. Damage to only the corpus callosum

can produce ‘anarchic hand’ syndrome in which the patient’s two hands fight to produce opposite effects – for example, one trying to undo a button while the other tries to do it up.

Single-cell recording in monkeys has explored the neuronal mechanisms of voluntary control of behaviour (Schultz, 1999), and newer methods of brain imaging have studied the functional anatomy of volition in humans. In an early study using PET, Chris Frith and colleagues (1991) compared conditions in which participants had either to repeat a given word or choose one. Subtracting the activity in one condition from the other revealed a difference in the left dorsolateral prefrontal cortex (DLPFC) and anterior cingulate. Other similar studies showed increased activity in DLPFC when actions were being selected and initiated. Reviewing such studies, Spence and Frith (1999) conclude that even the simplest motor procedures require complex and distributed neuronal activity, but the DLPFC seems to be uniquely associated with the subjective experience of deciding when and how to act.

Imagine you have to choose between your favourite brand of coffee or another slightly cheaper one, or between buying that expensive plane ticket now in case the flight sells out, or waiting in case it gets cheaper. Or perhaps, in an experiment, you can take 45 pence now or have a 50:50 chance of getting either a pound or nothing. What do you do? What does your brain do? These are the kinds of situations found in neuroeconomics, the study of the brain bases of economic behaviour (Politzer, 2008; Rangel, Camerer, and Montague, 2008; Glimcher and Fehr, 2013).

Even unconscious motivations can be measured. In one experiment, participants saw either a pound coin or a penny coin, and the force they exerted by gripping a handle determined how much they would get. They pressed harder for a share of the pound even when it was presented subliminally and they could not say which it was. Neuroimaging showed effects in part of the basal forebrain (Pessiglione et al., 2007). This suggests the perhaps slightly worrying idea that unnoticed stimuli around you are constantly affecting your motivations.

Now imagine that you are on a diet and tempted to succumb to a slice of chocolate cake. When your hand stops just in time, who or what stopped it? Decisions like this require what we think of as self-control to choose the option that is better in the long run over one which is immediately tempting. In one study of dieters making decisions about what to eat, fMRI scans suggested that the ventromedial prefrontal cortex was involved in encoding goal values, while activity in the DLPFC modulated these value signals when the participants were exercising self-control (Hare, Camerer, and Rangel, 2009).

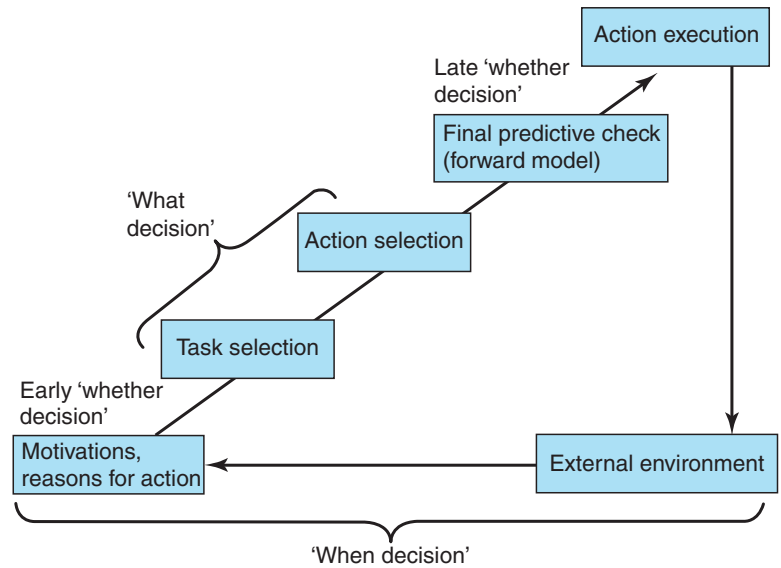


FIGURE 9.2 • ‘A naturalized model of human volition’ (Haggard, 2008, Figure 2). Volition is modelled as a set of decision processes that each specify details of an action. The decision whether to perform an action (‘whether-decision’) has both an early and a motivational component and a final predictive check. ‘What decisions’ specify which goal or task (from a range of tasks) to perform (‘task selection’) and the means by which to perform it (‘action selection’). The timing of voluntary actions often depends on the combination of environmental circumstances and internal motivations: an explicit ‘when decision’ is not always necessary (Haggard, 2008).

One way of investigating the experience of exercising conscious will is to ask what it feels like to have an urge to act. This may be in response to a sensory event like the tickle on your nose that you're longing to scratch, or a part of a pathology like Tourette's syndrome, which can involve swearing uncontrollably. Studies have investigated the neural correlates of both urges and suppression of urges, but it is hard to separate the urge from its inhibition, since prolonging the urge by definition means resisting it. You can also distinguish between the *what*, the *when*, and the *whether* aspects of intentional choice (Figure 9.2, and Brass et al., 2013).

Designing brain-imaging experiments that target only intentional action is difficult. In 'free-choice' scenarios with numerous trials, participants are usually instructed not to act stereotypically (e.g. not to simply alternate between responses) – that is, they are implicitly given a randomness instruction. In some studies (e.g. Soon et al., 2008), they are specifically asked not to make button selections in any kind of pattern. Parts of the fronto-parietal network might be involved in this strategic aspect of the tasks, by helping track the sequence of responses across trials in working memory (Lau et al., 2004a), even though this is not what is meant to be being tested. Another approach is to ask people to attend either to their intention to act, or to the action itself. One study (Lau et al., 2004b) found that attending to the intention led to stronger activation of the preSMA, supporting the 'shared circuit view' that the same areas are involved in objective control and the subjective experience of control. In contrast, however, in one of the rare experiments in which participants could choose between multiple options (by choosing what number to add to a systematic or unsystematic sequence), no overlap was found between areas thought to be involved in intentional choice and those correlated with participants' reports of feeling more freedom to choose (Filevich et al., 2013).

But the real problem for our purposes here is not just that isolating the neural correlates of 'free will itself' is fiendishly hard. It is the problem we keep coming up against in different guises: that having motivations and making decisions doesn't feel like neurons firing, whether in the SMA, DLPFC, or anywhere else. It feels as though there is something else – me, my own mind, my consciousness – that makes me free to act the way I want.

THE HALF-SECOND DELAY IN CONSCIOUSNESS

One of the most important questions to ask when it comes to the relationship between consciousness and voluntary action concerns timing. Does consciousness come early enough in the sequence of physical events that leads to an action to be able to exert a causal effect of its own?

To tackle this question, we need to go back to the late 1950s, when American neuroscientist Benjamin Libet began a series of experiments which led to the conclusion that about half a second of continuous neuronal activity is required for consciousness. This became popularly known as Libet's half-second delay (Nørretranders, 1998; McCrone, 1999). The issues raised are fascinating, and there have been many arguments over the interpretation of the results, so it is worth considering these studies in some detail.

In these early experiments, the sensory cortex of conscious, awake participants was directly stimulated with electrodes (Libet et al., 1979; Libet, 1982, 2004). They had all had invasive neurosurgical procedures carried out for therapeutic reasons and had given their informed consent. A small part of the skull was cut away, the somatosensory cortex exposed, and electrodes applied to stimulate it with trains of pulses which could be varied in frequency, duration, and intensity. The result, under certain conditions, was that the patients reported the definite conscious sensation of being touched on the skin of the hand, even though the only touch was a brief train of stimulation to the brain.

Using this method, Libet found a minimum intensity below which no sensation is elicited no matter how long the stimulation continues. But the surprising finding was that at this liminal intensity no experience was reported unless the stimulation continued for at least an average of 0.5 sec. At shorter durations, the intensity required to produce a reported experience rose very steeply. This length of time was roughly the same even when other variables, such as the frequency of pulses, were varied. The same was found in some subcortical pathways, but not in the dorsal columns of the spinal cord, on peripheral nerves, or on the skin.

Sensory stimuli normally produce an 'evoked potential' (an electrical potential recorded from electrodes on the scalp) in the relevant area of cortex as soon as 10–20 ms after presentation. Interestingly, Libet found that a single pulse applied to the thalamus or medial lemniscus (both part of the specific pathway leading to somatosensory cortex) could induce an evoked potential that appeared just the same as that induced by an actual sensory stimulus. But this single pulse never produced a conscious sensation, regardless of its intensity or the size of the evoked potential. Libet concluded that 'neuronal adequacy' for conscious sensation is only achieved after half a second of continuous stimulation in somatosensory cortex. Indeed, he suggested that 'it is sufficient duration *per se*, of appropriate neuronal activities, that gives rise to the emergent phenomenon of subjective experience' (Libet, 1982, p. 238). Obviously in ordinary life there is no direct stimulation of the cortex by electrodes, but the implication would be that a sensory stimulus (such as a touch on the skin) sets up continuing activity in somatosensory cortex and that this must continue for half a second if the touch is to be consciously perceived.

On the surface, this conclusion seems very strange. Does it mean that consciousness takes half a second to build up? And does this imply that our conscious perceptions lag half a second behind the events of the real world, far too late to consciously exert free will in many rapidly evolving situations?

Half a second is a very long time in brain terms. Signals travel along neurons at about 100 m per second, and can take less than a millisecond to cross a synapse. Auditory stimuli take about 8–10 ms to get from the ears to the brain and visual stimuli 20–40 ms. So a great deal can happen in half a second. This is true of behaviour as well. The reaction time to a simple stimulus (say pressing a button when a light comes on) can be as little as 200 ms, and recognising a stimulus takes more like 300–400 ms. Drivers can usually stop in response to a sudden danger in less than a second, and if we touch something dangerously hot our fingers will move out of the way in less than half a second. Could it really be that consciousness comes so much later?

Several further experiments tried to clarify what was going on.




FIGURE 9.3 • Diagram of Libet's experiment on subjective time order. A continuous stimulus train at 60 pulses per second was applied to sensory cortex (C), and a single pulse at threshold to the skin of the arm 200 msec later (S). The conscious experience of C (C-experience) was reported to occur approximately 500 msec after stimulation began, and was not reported at all unless stimulation continued for 500 msec. On this basis one might expect S-experience to occur 200 msec after C-experience. In fact it was reported to occur at approximately the time of the skin pulse, before the C-experience. These findings led Libet to propose the 'subjective referral of sensory experience backwards in time' (after Libet et al., 1979, Fig. 1).

It was already known that a strong stimulus to somatosensory cortex could interfere with sensations coming from a touch on the skin. So if consciousness really takes half a second to build up, then it should be possible to touch someone on the skin and then block the sensation by stimulating the cortex up to half a second later. This was exactly what Libet found. He stimulated the skin first and then the cortex. When the cortical stimulus came between 200 and 500 ms after the skin stimulus, the skin stimulus was not consciously felt. In other words, a touch on the skin that participants would otherwise have reported feeling was retroactively masked up to half a second later. This certainly seems to confirm the idea that neuronal adequacy for conscious perception is about half a second.

But how can this be? We do not experience things as happening half a second behind, and half a second is long enough that surely we would notice the delay. Libet checked this intuition by asking participants to report the subjective timing of two sensations. One was an ordinary stimulus to the skin; the other was a cortically induced sensation (the two feel noticeably different). The interval between them was systematically varied and participants had to say which came first. They consistently reported that the skin stimulus came first, even when it came almost at the end of the train of pulses. This is what might be expected from previous findings, but is also very strange. If half a second of neuronal activity is required for conscious perception, why is the skin stimulus (which must also be followed by half a second of appropriate activity to produce a conscious sensation) felt first?

Libet's controversial suggestion was that sensory experiences are subjectively referred back in time once neuronal adequacy has been achieved. In other words, what happens with any sensation is this. Information travels from, say, the skin, up to the relevant sensory area of cortex. If, and only if, activity continues there for the requisite half a second, the stimulus is consciously perceived. At that point it is subjectively referred back to the actual time at which it happened. If neuronal adequacy is not achieved (because the stimulus was not strong enough, because other brain processes suppressed the activity, or because a devious experimenter interfered directly in the cortex), nothing is consciously experienced.

How does subjective referral work? To what point in time is the experience referred, and how? Libet surmised that the primary evoked potential might act as a timing signal to which the sensation is referred back – or 'antedated'. Because evoked potentials occur so fast after peripheral stimulation, referring the sensation back to this point would mean no delay in conscious perception even though half a second of activity is required for neuronal adequacy. To test this, Libet and his colleagues (Libet et al., 1979) exploited two special features of what happens when the medial lemniscus (part of the pathway from the cutaneous receptors to

the thalamus) is stimulated. As with the cortex, long trains of pulses are required for neuronal adequacy, but unlike in the cortex, a primary evoked potential also occurs as it does when the skin is touched. The backwards referral hypothesis makes a clear prediction: that stimulation to medial lemniscus should be referred back in time to the start of its train of impulses, even though stimulation of cortex is not. In this final experiment, Libet again asked participants about the relative subjective timing of different stimuli. As predicted, he found that if a skin stimulus came at the same time as the start of a train of pulses to the medial lemniscus, the participants felt the two simultaneously – even though the train of pulses was only felt at all if the stimulation went on long enough to achieve neuronal adequacy.

What should we make of these findings? In general, critics have not agreed on serious weaknesses in the methods used or the specific results. The ideal way to be sure is to repeat the experiments, but medical advances mean that operations to expose the brain are now very rare. So the experiments are unlikely ever to be replicated. We are probably best, then, to assume that the findings are valid. The real controversy surrounds how to interpret them.

Libet's own interpretation is his 'time-on theory' of consciousness. This has two components: first, that consciousness can occur only when neural activity continues long enough for neuronal adequacy (usually about 500 ms), and second, that activity with a shorter duration can still be involved in an unconscious process or converted into a conscious one by increasing its duration. He suggests that attention may work by increasing the excitability of certain areas so as to lengthen the duration of activity and so achieve the time-on for consciousness (Libet, 2004). On his view, unconscious processes really do 'become conscious' when neuronal adequacy is achieved. He says that 'when the duration of repetitive similar activations of appropriate neurons reaches a certain value, then the phenomenon of awareness emerges' (2004, pp. 58–59).

This theory provides an answer to the question we focused on in the last chapter: what is the difference between conscious and unconscious processes? According to Libet, the difference is whether neuronal adequacy is reached or not. To compare this with one contrasting example, when Milner and Goodale suggest that processing for perception in the ventral stream leads to consciousness, while dorsal stream processing for action does not, Libet (1991) argues that the important difference for consciousness is not the brain areas where the processing occurs, nor what kind of activity it is, nor what it leads to, but only whether it continues for long enough.

Libet also makes some much more controversial suggestions. In particular, he claims that the evidence for backwards referral raises problems for materialism and the theory of psychoneural identity (i.e. that consciousness and neural activity are the same thing). He even considers 'the possibility that physical events are susceptible to an external "mental force" at the micro level, in a way that would not be observable or detectable' (Libet, 2004, p. 154). Roger Penrose (1994a and 1994b) also believes that the phenomena uncovered in these experiments challenge ordinary explanations and demand reference to nonlocality and quantum theory. Similarly, Karl Popper and John Eccles claim that 'This antedating procedure does not seem to be explicable by any neurophysiological process. Presumably it is a strategy that has been learnt by the self-conscious mind' (1977, p. 364). In

DID MY THOUGHTS CAUSE THIS ACTION?

'when the duration [...] reaches a certain value, then the phenomenon of awareness emerges'

(Libet, 2004, pp. 58–59)

other words, they think that intervention by the nonphysical mind is required to explain subjective antedating. On their view, Libet's results provide evidence for dualism – a claim that others firmly reject (e.g. Churchland, 1981; Dennett, 1991).

Despite his belief in a mental force, Libet himself points out that subjective referral in space has long been recognised and so we should not be surprised to find subjective referral in time as well. Although it may seem odd that we experience objects as 'out there' when vision depends on our brain 'in here', this kind of projection is not magical, he says – and nor is subjective referral. Given the widely dispersed activity in the central nervous system, we should expect a mechanism that coordinates subjective timings. Subjective referral to the evoked potentials does just that.

Let's return to our scenario from Chapter 7, about turning to look at the person coming into the room. Which comes first, the movement to see who it is, or the awareness? If Libet is right, then conscious perception of the noise cannot occur unless there is at least half a second of continuous neural activity after the noise begins. Since we often react far faster than that, this means that the causal sequence cannot be 1) consciously hear sound, and 2) turn round to look.

The previous paragraph was carefully worded. It said that conscious perception cannot occur *unless* there is at least half a second of continuous neural activity after the noise begins – which is indeed suggested by Libet's results. What is not necessarily implied, though it is often assumed, is something like this: after the noise occurs there is a lot of unconscious processing. Then, after half a second, the noise 'becomes conscious' or comes 'into consciousness'. At that point, it is antedated so that it seems to have occurred earlier, at the right time. On this view, consciousness really does trail along half a second behind the events of the real world, but we don't realise it.

The difference between these two descriptions is important. The first doesn't commit to a time at which consciousness happens or emerges. The second does: it assumes there is a fact of the matter about when processing becomes conscious. In Chapter 6, we considered the distinction between clock time and perceived time, and reviewed reasons to question the very idea that there is a measurable time at which subjective experience happens. This means we should also question Libet's view that the *experience itself* can be timed, and that consciousness happens when neuronal adequacy is achieved. In any case, the findings from Libet's experiments on the timing of consciousness remind us to pay careful attention to timing when we ask whether and how consciousness contributes to 'freely willed' action.

THE ROLE OF CONSCIOUS WILL IN VOLUNTARY ACTION

Hold out your hand in front of you. Now, whenever you feel like it, consciously, deliberately, and of your own free will, flex your wrist. Keep doing this for some time – until your arm gets too tired. Just flex your wrist whenever you want, and try to observe what goes through your mind as you do so. If you don't want to do it at all that's fine – that is your conscious decision. If you want to do it frequently, that is fine too. Now ask yourself what started the movement, or prevented it, each time. What caused your action?

This simple task formed the basis of one of the best-known experiments in the history of consciousness studies: Libet's (1985) study of 'Unconscious cerebral initiative and the role of conscious will in voluntary action'. Since the 1960s, it had been known that voluntary motor actions are preceded by a 'readiness potential' (RP): a slow negative shift in electrical potential that can be recorded from electrodes on the scalp up to a second or more before the action takes place. This long interval (on average 800 ms) prompted Libet to wonder 'whether the conscious awareness of the voluntary urge to act likewise appears so far in advance' (1985, p. 529).

He reasoned that if a conscious intention or decision initiates the action, then the subjective experience of intending should come first, or at the latest together with the start of the cerebral processes. This was what his experiment investigated. He needed to time three events: the start of the action itself, the start of the RP, and the moment of the conscious decision to act, which he called 'W' for 'will'.

Timing the action itself was easily done with electrodes on the appropriate muscle. Timing the RP was also relatively straightforward, although the change in potential that marks the beginning of the RP can be clearly seen only when averaged over many repetitions (note that the RP is a gradual slope, not a sudden change). Libet had his five participants (plus one other whose data were mostly not usable) carry out the flexion forty times in each series of trials. Using the time of the action as a reference, these forty trials could then be averaged even though the participant freely chose when to move (Libet et al., 1983). The source of this RP was thought to be the SMA.

The real problem is how to measure W, the moment when someone becomes conscious of the urge, or will, to move. If you ask people to say 'Now' when they feel like moving, the action of speaking may not only interfere with the wrist movement, but may also involve its own RP and another delay. So Libet used the following method. A spot of light revolved once every 2.56 seconds, in a clock circle on a screen showing twelve segments of 5. Participants were asked to watch the spot carefully and then, after they had flexed their wrist, to report where the spot was at the moment they felt the urge to move. So on each trial they would flex their wrist and then say



ACTIVITY 9.2

Libet's voluntary act

Libet's experiment is complex, and the arguments about its interpretation are fierce. It will help you understand them if you practise the role of one of his participants. Having tried it, you will be much more likely to think up, for yourself, all the classic objections to Libet's conclusion.

As a class demonstration, ask everyone to hold out their right arm in front of them and then, whenever they feel like it, consciously, deliberately, and of their own free will, flex their fingers or wrist. They should perhaps do this about forty times (as in Libet's experiment), but since people vary in speed (and some may freely choose not to do it at all), about two minutes is usually enough.

Now ask your participants what they think started the action. They might suggest it was their inner self, or a thought, intention, or feeling that started it, or that a stream of brain events was responsible. Ask whether the action seemed free or not. Could they have done otherwise? Is this a good model for a 'spontaneous voluntary act'?

Now you need to time 'W': the time at which they decided to act. Stand in front of the group with your arm straight out, and use your own hand to represent the rotating light spot (for a large audience, hold a bright object in your hand to make it more visible). Make sure your hand rotates steadily clockwise from the viewers' point of view at roughly one revolution every two seconds (Libet's spot went a little slower but 1 in 2 works well; practise first). Now ask the audience to do the same flexing task as before, but this time they must, after they have acted, shout out the clock position (from 1 to 12) at the moment when they decided to act. You now have a room full of people shouting out different times all at once. One question is, can we easily do this? Most people find they can.

Libet measured three things: the start of the action itself, the start of brain activity leading to the action, and the decision to act. Ask yourself which you expect to come first, or get everyone to put up their hands.

You are now ready to discuss Libet's experiment and what his results really mean.




FIGURE 9.4 • In his experiments on voluntary action Libet (1985) timed three things: M, the movement of the hand or wrist; RP, the readiness potential detected from motor cortex using EEG; and W or ‘will’. W was timed by asking participants to watch a revolving spot and say (afterwards) where the spot was when they decided to move.

‘If the brain can initiate a voluntary act before the appearance of conscious intention [...] is there any role for the conscious function?’

(Libet, 1985, p. 536)

‘15’ or ‘35’, meaning that this was the position of the spot at the moment they decided to act. A control series, in which participants reported the time of a skin stimulus by using the clock method, showed that their estimates were generally accurate and slightly in advance of the actual stimulus. In another control, participants were asked to time their awareness of actually moving (M). They had no trouble following these instructions nor in discriminating M from W – W being, on average, 120 ms before M. Using these controls, Libet was convinced that the timing of W was sufficiently accurate. He could now answer his question: which comes first, the start of the readiness potential or the conscious decision to act?

The answer was clear. The RP came first. On average RP started 550 ms (+/- 150 ms) before the action and W only 200 ms before. In the debriefing conducted after each series of forty trials, participants said that on some trials they had been thinking about the action some time in advance, or preplanning

it. On these trials the RP began over a second before the action, but for series in which all forty acts were reported as fully spontaneous, the RP began 535 ms before the action, and W just 190 ms before the action. Further analysis showed that this held for different ways of measuring both RP and W. In conclusion, the conscious decision to act occurred approximately 350 ms *after* the beginning of RP.

What should we make of this finding? With Libet, we may wonder: ‘If the brain can initiate a voluntary act before the appearance of conscious intention [...] is there any role for the conscious function?’ (Libet, 1985, p. 536). That is the crux. These results seem to show that consciousness comes too late to be the cause of the action.

For those who accept the validity of the method, there are two main ways of responding to Libet’s results. The first is to say, ‘Well, that’s obvious! If consciousness came first, it would be magic’. Presumably this ought to be the standard reaction of anyone who denies dualism. Indeed, the result should have been completely unsurprising. Instead, even though most psychologists and philosophers deny being dualists or believing in magic, these results caused a furore. Not only was there a wide-ranging debate in *Behavioral and Brain Sciences*, but the experiment continued to be frequently cited, and hotly argued over, more than twenty years later (Libet, 1999, 2004).

The second response is to seek some remaining causal role for consciousness in voluntary action. Libet took this route and argued as follows. It is possible to believe, he said, that conscious intervention does not exist and the subjective experience of conscious control is an illusion, but such a belief is ‘less attractive than a theory that accepts or accommodates the phenomenal fact’ (i.e. the fact about how it feels), and is not required even by monist materialists (Libet, 1999, p. 56). For example, Roger Sperry’s emergent consciousness is a monist theory in which consciousness has real effects. For Sperry, mental activity emerges from neural activity and can then have effects back on it. By limiting these effects to ‘supervening’, not ‘intervening’, he could remain a determinist (though Libet notes

that in the end he abandoned determinism [pp. 168–169]). The results are also compatible with dualist interactionism (Popper and Eccles, 1977) or with ‘the possibility that physical events are susceptible to an external “mental force” at the micro level’ (Libet, 2004, p. 154). Libet therefore proposed

that conscious control can be exerted before the final motor outflow to select or control volitional outcome. The volitional process, initiated unconsciously, can either be consciously permitted to proceed to consummation in the motor act or be consciously ‘vetoed’.

(Libet, 1985, p. 536–537)

The idea, then, is that unconscious brain events start the process of a voluntary act but then, just before it is actually carried out, consciousness may say either ‘yes’ or ‘no’: the action either goes ahead or not. This would happen in the last 150 to 200 ms before the action. Libet provides two kinds of evidence for this conscious veto. First, participants sometimes reported that they had an urge to act but then aborted or suppressed the action before it happened. Unfortunately, the neural correlates of aborted self-timed actions cannot be measured because averaging over many trials needs a movement signal to act as the time cue. So in additional experiments, participants were asked to move at pre-arranged times, and then abort some of the actions, allowing the averaging to be done. These showed ramplike pre-event potentials that then flattened or reversed about 150–250 ms before the preset time. This suggested to Libet that the conscious veto interfered with the final development of the RP.

In this way, Libet was able to retain a causal role for consciousness in voluntary action. He concluded that his results are not antagonistic to free will but rather illuminate how free will operates. When it comes to morality and matters of conscience, we can still be expected to behave well. Although we cannot consciously control having an impulse to carry out an unacceptable action (whether rape or murder or stealing sweets in the supermarket), we can be held responsible for consciously allowing its consummation – or not. As Richard Gregory characteristically punned it, ‘We don’t have free will, but we do have free won’t’ (1990). The idea has since acquired more support from a study finding a similar ‘point of no return’ about 200 ms before movement onset: before this point, the movement can still be vetoed (Schultze-Kraft et al., 2015).

As with Libet’s earlier experiments, the debate following publication of his results raised both philosophical and methodological problems (undated references in this section all refer to commentaries following Libet [1985]). Eccles used the data to support his dualist-interactionist theory, and David Rosenthal (2008) has argued that the findings are just what a HOT theory of consciousness would predict: a mental state is conscious only if it is the object of a higher-order mental state, which you would expect to come after the decision itself. So he says

Self-initiated act: sequence




FIGURE 9.5 • According to Libet the sequence of events in a self-initiated voluntary act is as shown. Preplanning (RPI) occurs as much as a second before the movement. For spontaneous actions without preplanning, activity (RPII) begins about half a second before the movement. Subjective awareness of the will to move appears about 200 msec before the movement. Subjective timings of a randomly delivered skin stimulus (S) averaged about –50 msec from actual time (Libet, 1999, p. 51).

*'We don't have free will,
but we do have free
won't'*

(Gregory, 1990)

that the RP is the volition which first initiates the action and only then becomes conscious; the actions and the consciousness are caused by the RP, rather than consciousness causing the action. Others, though, have criticised Libet for his unstated dualist assumptions (Wood), and even for 'double dualism' (Nelson) and 'metaphysical hysteria' (Danto). These criticisms revolve around the way that Libet compares physical with mental events, and tries to defend what seems to be a magical 'conscious control function' in his proposed veto.

Other criticisms turn on the question of whether the RP is best thought of as the neural basis of the urge to act – that is, as one of the complex motivational features that contribute to agency – rather than as the neural basis of the decision to flex my wrist now: as one tributary not *the origin* (Bayne, 2011). Tim Bayne also invites us to think more carefully about what the intuitive or 'folk' concept of free will challenged by Libet's work really involves. Does 'free' will require the initiating conscious decision to be an 'uncaused cause' (to have no causal chain stretching

back beyond it), or would 'freedom' be compatible with the idea that as long as a conscious decision is the immediate cause, that decision can itself be caused by any number of actions based on previous conscious decisions? In that case, though, how far do we have to trace back before we find consciously exercised freedom, and what will it look like?

The main methodological criticisms concerned the nature of the task and the method of timing *W*. Several commentators argued that the task was not a good model of volition in general. This was partly because the action was so trivial, and partly because the participants could choose only the timing of their action, not the act itself, so any conscious willing would have happened before their decision about *when* to act. The results should not, therefore, be generalised to other, more complex willed actions, let alone to questions of moral responsibility (Breitmeyer, Bridgeman, Danto, Näätänen, Ringo).

Psychologist Richard Latta raises questions about backwards referral. If perception of the position of the spot *and W* are both subjectively referred backwards in time, then the two will be in synchrony, but if *W* is not referred back, then the timing procedure is invalidated. In response, Libet points out that backwards referral is not expected for the spot because the time at which the participants became aware of the spot was not the issue, only its position when they felt the urge to act. If this still seems obscure, we might imagine participants who had the experience of deciding to move exactly as the spot reached 30. It would not matter how long this perception of



CONCEPT 9.1

VOLITION AND TIMING

Why don't you laugh when you tickle yourself? In experiments using a robot tickling arm, fMRI showed activity in secondary somatosensory cortex, anterior cingulate, and cerebellum. This was reduced when participants tried to tickle themselves with the arm, and timing proved critical. When the self-tickling sensations from the robot arm were delayed by more than 200 msec, the sensation became ticklish again (Blakemore et al., 1999).

Timing is critical to the experience of will in other ways too. Wegner and Wheatley's (1999) experiments on the 'priority principle' suggest that the timing of an event can affect whether we feel we willed it or not. Could the opposite also be true, and the perceived time of an event depend on its cause? Although this may seem peculiar, there is evidence that 'when we perceive our actions to cause an event, it seems to occur earlier than if we did not cause it' (Eagleman and Holcombe, 2002).

simultaneity took to become conscious because they could report this spot position at their leisure.

The whole method of timing W was also criticised, as was the adequacy of using a skin stimulus as a control to test the accuracy of the timing, and the failure to allow for delays involved in each, or in switching attention between the spot and W (Breitmeyer, Rollman, Underwood, Niemi). There have also been proposals that instead of reflecting preconscious motor preparation, the readiness potential might instead be the result of an averaging of random noise that exceeds a certain threshold – that is, it might not tell us anything about readiness for a specific action (Schurger, Sitt, and Dehaene, 2012).

Some of these criticisms are undermined by subsequent replications. For example, British psychologist Patrick Haggard and his colleagues not only replicated the basic findings, but showed that awareness of one's own actions is associated with a premotor event (lateralised RP) after the initial intention and preparation, but before the motor command is sent out (Haggard, Newman, and Magno, 1999). Comparing trials with early and late awareness, they found that the time of awareness covaried not with the RP but with the lateralised RP, concluding that 'the processes underlying the LRP may cause our awareness of movement initiation' (Haggard and Eimer, 1999, p. 128). Haggard and Libet (2001) then debated the implications of these results.

A 2008 study by Chun Siong Soon and colleagues in Leipzig updated Libet's experiment using fMRI, tweaking the design to try to circumvent some of the criticisms, particularly with regard to the timing of W. Instead of using a clock face for timing, participants were presented with consonants in the middle of a screen, one at a time for 500 ms each, and asked to passively observe the stream of letters. This made the sequence unpredictable, unlike the hand moving round the clock, to avoid them anticipating their decision or choosing in advance a time at which to move. They were told to relax (not to be too eager to press a button when the letters first appeared, nor to maintain a constant state of readiness to move), and to press either the left or the right button with the index finger of the corresponding hand as soon as they felt the urge to do so. They were asked to remember the letter that was on the screen when they decided which button to press, not when they

In experiments on voluntary action and conscious awareness, Patrick Haggard and colleagues at University College London used Libet's clock method for participants to time the onset of four single events: a voluntary key press, a muscle twitch produced by stimulating their motor cortex with transcranial magnetic stimulation (TMS), a click made to sound like TMS, and a tone. Next, in the voluntary condition, they pressed a key and a tone sounded 250 msec later. In the TMS condition, their finger twitched involuntarily and the tone followed, and in a control condition just the click was used. In each case, they reported the time of the first event and when they heard the tone.

In this second stage, large perceptual shifts were found as compared with the single-event case. The voluntary key press and the time of the tone were reported as being closer together, whereas the involuntary twitches (caused by TMS) and the tone seemed further apart. There was no effect for sham TMS and the effect was greatest for shorter time intervals. The effect is known as 'intentional binding' and its strength can be affected by predictability, feedback, and beliefs (Moore and Obhi, 2012).

What does this imply for consciousness? The experimenters themselves claimed that

the perceived time of intentional actions and of their sensory consequences [...] were attracted together in conscious awareness, so that subjects perceived voluntary movements as occurring later and their sensory consequences as occurring earlier than they actually did.

(Haggard, Clark, and Kalogeras, 2002, p. 382)

This interpretation is a form of Cartesian materialism, implying that events are perceived and manipulated 'in conscious awareness'. A more sceptical interpretation is that the important processes of timing and discriminating between self-caused and external events happen without anything being 'in' or 'out' of consciousness.

Haggard asks whether the conscious experience of owning an action depends on predicting the coming action or inferring agency afterwards. From these and further experiments on timing, he concludes that 'The phenomenology

of intentional action requires an appropriate predictive link between intentions and effects, rather than a retrospective inference that ‘I caused the effect’ (Haggard and Clark, 2003, p. 695).

actually pressed it, and to choose it from a selection of three that appeared after the button press.

Soon and colleagues found that by studying activity in the prefrontal and parietal cortex, they could predict the outcome of a left-or-right decision up to ten seconds before participants themselves became aware of their choices. They conclude that ‘This delay presumably reflects the operation of a network of

high-level control areas that begin to prepare an upcoming decision long before it enters awareness’ (p. 543). The prediction rate, however, is relatively low (only slightly above chance). This raises the question of whether these early predictive cues are the precursors of intention or rather biasing signals that influence the choice process without constituting it (see also Haynes, 2011). Interestingly, two of the brain areas that carry early predictive information are the central nodes of

the default mode network (see Chapter 8), which means that in the context of decision-making they might play an evaluative role, providing background information that biases decision-making (Brass et al., 2013, p. 7) long before the moment when we become aware of making a choice.

But what is this moment? What is the ‘time of awareness’, or the time when a decision ‘enters awareness’? The most radical critique is given by Dennett, who asks us to join him in the following ‘all-too-natural vision’ of Libet’s wrist-flexing task (Dennett, 1991, p. 165).

Unconscious intentions start out somewhere deep in the brain and then, gradually becoming more definite and

powerful, make their way up to where ‘I’ am. At this point, they ‘enter consciousness’ and ‘I’ have the experience of deciding to act. Meanwhile, representations of spots on a clock face have been streaming up from the retina, gradually becoming more definite in brightness and location, until they too reach consciousness and ‘I’ can watch them parading past. So at the very moment when the intention appears in consciousness, ‘I’ can say where the spot was.

As Dennett points out, this is so easy to visualise. Isn’t that how it has to be when two things happen together in consciousness? No. Indeed, he says it cannot be. There is no place or system in the brain where all the things currently ‘in consciousness’ meet together, there is no time at which things ‘enter consciousness’, and there is no self watching the display in that non-existent place. To try to




FIGURE 9.6 • Haggard and colleagues (2002) report that the judged time of a tone changes as a function of the delay between the tone and a previously executed voluntary act. As the delay is lengthened (a–c), the time misestimation is reduced. Mean judged time is represented by thought bubbles. In the experiment, time judgements are always retrospective, which is why they can appear to precede the actual times of occurrence on the timelines (Eggleman and Holcombe, 2002).

escape this impossible vision, some theories hold that consciousness is a matter not of arriving at a place, but of exceeding a threshold of activation in a distributed system or network. So things can 'enter consciousness' while staying put. This changes the imagery, but not the basic mistake, says Dennett. In this version, there has to be some moment at which physical activity achieves the special state, and some way in which it acquires the special quality of *subjectivity*, so becoming 'my conscious decision'. This moment is what is timed in Libet's experiment.

These two visions may sound different, but they both entail a Cartesian theatre: a 'headquarters' – whether centralised or distributed – in which different things 'come together' in consciousness, and from which consciousness does its controlling. Only with such a vision can you imagine, as Libet does, that 'the conscious function' can trigger some actions and veto others. In this way, says Dennett, both Libet and most of his critics remain trapped in the Cartesian theatre.

One way out is to abandon the notion that there is an answer to the question 'what is in my consciousness now?' You can retain the idea that the brain makes judgements of simultaneity – and often very accurate ones – but only because brain mechanisms time events and produce behaviours or statements based on those timings. There is no additional 'you' with a privileged view of the contents of your consciousness and the conscious power to act.

So does Dennett believe that free will is illusion? He says not (Dennett, 2003), but his reasons may cause some confusion because his view neatly fits the definition of 'illusion' we are using here: that an illusion is something which is not as it seems. He explains that if you believe that free will springs from an immaterial soul shooting arrows of decision into your brain, then there is no free will at all, but if you believe that free will might be morally important without being supernatural, then 'free will is indeed real, but just not quite what you probably thought it was' (p. 223). Human freedom is not magic, but an evolved capacity for weighing up options and dealing with multiple choices. We are then left with the question of whether it makes sense to carry on using the term 'free will' to refer to something so unlike the freedom most people imagine when they say it. We may also find ourselves asking why we attribute freedom to the will rather than to the person doing the willing – another instance of the mereological fallacy in action, perhaps. But these questions take us into a whole different realm of the philosophy and psychology of language use, which is beyond our scope here.

So where does all this get us? If personal conscious will is a real force acting on the brain, as James, Libet, Eccles, and others would have it, then there is no mystery about why we *feel* as though we can consciously exert free will. We can. On the other hand, if free will is an illusion, then we have a new mystery. Why do we *feel* as though our conscious decisions cause our actions when they do not?

To find out, we must ask about the origins of the *experience* of will, asking not whether free will exists, but what creates the feeling of exerting our will and what makes that feeling also feel 'free'. There are many overlapping concepts here: agency, control, volition, will, and freedom, to name the most common. We will try to be faithful to the different terms researchers use, but you will have to make your own mind up about whether they are all investigating the same thing, or whether there is even a unitary thing to be investigated.

'free will is indeed real, but just not quite what you probably thought it was'

(Dennett, 2003, p. 223)

THE EXPERIENCE OF WILL

THE ILLUSION OF NO WILL

In 1853 the new craze of spiritualism was spreading rapidly from the United States to Europe (see Chapter 15). Mediums claimed that spirits of the dead, acting through them, could convey messages and move tables. Appreciating the challenge to science, and infuriated by public hysteria, the famous physicist and chemist Michael Faraday (1853) investigated what was going on.

In a typical table-turning séance, several sitters sat around a table with their hands resting on the top. Although they claimed only ever to press down, and not sideways, the table would move about and spell out answers to questions. They all said that the table moved their hands, not that their hands moved the table. In an ingenious experiment, Faraday stuck pieces of card between the sitters' hands and the table top, using a specially prepared cement that allowed the cards to move a little. Afterwards he could see whether the card had lagged

behind the table – showing that the table had moved first as the sitters claimed – or had moved ahead of the table. The answer was clear. The card moved ahead, so the force came from the sitters' hands. In further experiments, Faraday fixed up a visible pointer which revealed any hand movements. When the sitters watched the pointer, 'all effects of table-turning cease, even though the parties persevere, earnestly desiring motion, till they become weary and worn out' (Faraday, 1853, p. 802). Visual feedback sensitised the sitters to their muscular activity in a way that proprioceptive feedback had not been able to. He concluded that unconscious muscular action was the only force involved.

Psychologist and magician Jay Olson (Olson et al., 2016) explored a twenty-first-century version of the paranormal in his 'simulated thought insertion' study. For the 'Mind-Reading Task', participants lay in a dummy brain scanner and were told that the machine was part of a 'Neural Activation Mapping Project', and could read and influence their thoughts. The scanner made realistic noises, and a printer in the next room supposedly printed out (along with lots of technical-looking but meaningless statistics) the number they were thinking of, but with occasional mistakes to make it seem more realistic. Participants were convinced, and expressed surprise, amusement, confusion, or discomfort at the idea of the machine reading their thoughts.



FIGURE 9.7 • A spiritualist séance from 1853. In table turning, or table tipping, the sitters believed that spirits moved the table and that their own hands just followed. Faraday proved that the movements were due to unconscious muscular action.

Next, for the 'Mind-Influencing Task', they were told that the machine would randomly choose a number and try to influence them to select it, by manipulating 'natural electromagnetic fluctuations in the brain'. Participants again believed in the machine's powers, some reporting having a hot face or feeling a pulsation when the machine was influencing them, and others referring to an unknown

source directing them towards specific numbers. When later asked whether they could guess something they had not been told about the experiment, only nine out of sixty expressed some suspicion. In this task, participants gave higher ratings for involuntariness, and took longer to choose their numbers, than in the 'Mind-Reading Task'. In interviews, they spoke of how in the mind-influencing condition the decision 'just happened' or the number 'came out of nowhere. So I felt like it . . . wasn't my choice' (Olson et al., 2016, p. 21). Some mentioned trying to change the number and feeling they couldn't, whether it was their own brain being disobedient ('my brain just told me no, that's not the number') or the power of the machine dictating to them ('once the magnet turned on . . . I got 4'), or a voice or force or image trying to distract them.

Hints of a similar effect – causing something to happen without feeling responsible – were found decades earlier in the 'precognitive carousel' (Dennett, 1991). In 1963 the British neurosurgeon William Grey Walter tested patients who had electrodes implanted in their motor cortex as part of their treatment. They sat in front of a carousel slide projector and could press a button, whenever they liked, to see the next slide. Unbeknown to them, the slide was advanced not by the button-press but by amplified activity from their own motor cortex. The patients were startled, saying that just as they were about to press the button, the slide changed all by itself. When pressing the button, they also found themselves worrying about accidentally changing the slide twice. Perhaps with a longer delay between the cortical activation and the change of slide they would have noticed nothing amiss, but sadly Grey Walter did not experiment with variable delays. Nevertheless, without relying on the kind of artificial judgement about the timing of will required in Libet's experiment, this simple finding of surprise demonstrates that under certain conditions we can actually be in control of our actions without *feeling* that we are.

A similar mismatch occurs as a symptom of schizophrenia (Mullins and Spence, 2003). Many people with schizophrenia believe that their actions are controlled by aliens, by unspecified creatures, or even by people they know. Others feel that their own thoughts are controlled by evil forces, or inserted into their minds. This disconnection between voluntary action and the *feeling* of volition is deeply disturbing.

THE ILLUSION OF WILL

Can it happen the other way around? Can we have the sense of willing an action for which we are *not* responsible? Magicians have long made observers believe they have freely chosen a card or number, when in fact it was forced. Other experiments by Olson and others show how easy it is to influence people's choices without them noticing, even if this involves the magician actively handling the card in question (Shalom et al., 2013; Olson et al., 2015). The outcomes of our actions, and how likely those outcomes are, can also affect how responsible we feel for them: a heightened sense of agency may result from 'nice surprises', where an action outcome is both positive and unexpected, without there being a symmetrically reduced sense of agency for nasty surprises. On the other hand, an anticipated sense of agency is lost when an outcome is predictably positive or (even more so) negative: 'affective context may change the experience of *the nature and the quality of the act*' (Christensen et al., 2016, p. 8). This has important consequences



PROFILE 9.1

Daniel Wegner (1948–2013)



Having started a degree in physics, Daniel Wegner changed to psychology as an anti-war statement in 1969, and became fascinated with questions of self-control, agency, and free will. He did numerous experiments on how the illusion of free will is created, and on the effects of trying *not* to think about something. 'Try not to think about a white bear', he suggests.

From the age of 14, Wegner helped his mother, a piano teacher, run her music studio and taught piano twice a week after school. He not only played the piano but had four synthesisers for composing techno. When Professor of Psychology at Harvard University, he started all his classes with music. A colleague called him 'one of the funniest human beings on two legs'. He enjoyed studying 'mindbugs', those foibles of the mind that provide fundamental insights into how it works, and believed that conscious will is an illusion.

'Our sense of being a conscious agent who does things comes at a cost of being technically wrong all the time'

(Wegner, 2002, p. 342)

DID MY THOUGHTS CAUSE THIS ACTION?

for how we understand legal responsibility: 'Reduced responsibility could correspond to a fact of human psychology, rather than a hopeful story to avoid punishment' (p. 9).

These examples reveal, from both directions, the important difference between actually causing something to happen and having the *feeling* of causing it. As Daniel Wegner puts it, 'The feeling of doing is how it seems, not what it is' (2002, p. 342), and he has examined in detail the mechanisms that produce this *experience* of conscious will.

Imagine that you are standing in front of a mirror with screens arranged so that what look like your arms are actually someone else's. In your ears you hear instructions to move your hands, and just afterwards the hands carry out those same actions. Experiments showed that, in such a situation, people felt they had willed the movements themselves.

This is 'the mind's best trick', says Wegner (2003). Does consciousness cause action? A lifetime of experiences leads us to believe so, but in fact experiences of conscious will are like other judgements of causality, and we can get the judgement wrong. Indeed, his stark conclusion is that 'Our sense of being a conscious agent who does things comes at a cost of being technically wrong all the time' (2002, p. 342). American psychologist Sam Harris agrees: 'There is no question that our attribution of agency can be gravely in error. I am arguing that it always is' (2012, p. 25).

Wegner proposes that 'The experience of willing an act arises from interpreting one's thought as the cause of the act' (Wegner and Wheatley, 1999, p. 480), and that free will is an illusion

created in three steps. First, our brain sets about planning actions and carrying them out. Second, although we are ignorant of the underlying mechanisms, we become aware of thinking about the action and call this an intention. Finally, the action occurs after the intention, and so we leap – erroneously – to the conclusion that our intention caused the action.

This is similar to James's theory of deliberate actions, proposed over a century earlier. First, various reinforcing or inhibiting ideas compete with each other to prompt a physical action – or not. Once one or the other finally wins, we say we have decided. 'The reinforcing and inhibiting ideas meanwhile are termed the *reasons* or *motives* by which the decision is brought about' (1890, ii, p. 528). Note that both these theories explain how the powerful *feeling* that we willed an action might come about, whether or not we have free will. Interestingly, James and Wegner come to opposite opinions on this central question.

Wegner suggests that there are three requirements for creating the experience of willing: the thought must occur before the action, the thought must be consistent with the action, and the action must not be accompanied by other plausible causes. To test these proposals, Wegner and Wheatley (1999) carried out an experiment inspired by the traditional ouija board, which, like Faraday's turning tables, depends on unconscious muscular action. The ouija board (the name comes from

the French and German for 'yes') is used to try to contact spirits. Several people place their fingers on an upturned glass in the middle of a ring of letters and the glass then moves, spelling out words. Like Faraday's spiritualist sitters, people are usually sure they did not push the glass. But in Wegner's version, participants were explicitly instructed to exert control over the board's movements. The glass was replaced by a 20 cm square board mounted on a computer mouse, and this moved a cursor over a screen showing about fifty small objects. Fifty-one undergraduates were tested and each was, unbeknown to them, paired with a confederate. We'll call the two people Dan (the participant) and Jane (the confederate).

Dan and Jane sat facing each other across a small table and were asked to place their fingers on the little board and to circle the cursor over the objects. They were asked to stop every 30 seconds or so, and then rate how strongly they had intended to make that particular stop. Each trial would consist of 30 seconds of movement, during which they might hear words through headphones, and 10 seconds of music, during which they were to stop. Dan was led to believe that Jane was receiving different words from his, but actually she heard instructions to make particular movements.

On four trials she was told to stop on a particular object (e.g. swan) in the middle of Dan's music. Meanwhile Dan heard the word 'swan' 30 seconds before, 5 seconds before, 1 second before, or 1 second after Jane stopped on the swan. In all other trials the stops were not forced. The results confirmed what Wegner calls the 'priority principle': that effects are experienced as willed when the relevant thoughts occur just before them. On forced trials, participants gave the highest rating for 'I intended to make the stop' when the word came 1 or 5 seconds before the stop, and the lowest when it occurred 30 seconds before or 1 second after.

Wegner's principles might underpin the illogical feeling many people have that they can magically influence events they care about. In further studies, he and his colleagues gave people the impression that they had harmed someone else through a voodoo hex (Pronin et al., 2006). The effect was stronger among those who had first been induced to harbour evil thoughts about their victim. During sports events, people often superstitiously wear team kit, or urge their favourite player to run a bit faster or score the crucial goal, even if they are watching on TV and their encouragements can make no difference. In studies of baseball shooting, observers were more likely to think they had influenced a friend's success if they had first visualised success (Pronin et al., 2006). In these ways, the mechanisms that give rise to the feeling of willing can even extend to 'everyday magical powers' that we know are impossible. We may well do these things to feel more involved and less helpless as our team battles it out. Nonetheless, we often experience the sense that maybe, just maybe, really wanting something to happen could make a difference. Perhaps this impression is equally mistaken when it applies to our own actions. 'Believing that our conscious thoughts cause our actions is an error based on the illusory experience of will – much like believing that a rabbit has indeed popped out of an empty hat' (Wegner and Wheatley, 1999, p. 490). For Wegner, the illusion of will really is like magic and arises for the same reason. Yet, once again, we must remember that an illusion is not something that does not exist, and illusions can have powerful effects. Wegner concludes:

The fact is, it seems to each of us that we have conscious will. It seems we have selves. It seems we have minds. It seems we are agents. It seems we

*'Compatibilism is
'wretched subterfuge
[...] petty word-jugglery'*

(Kant, 1788/1956,
pp. 189–190)

• SECTION THREE: BODY AND WORLD

cause what we do. Although it is sobering and ultimately accurate to call all this an illusion, it is a mistake to conclude that the illusory is trivial.

(Wegner, 2002, p. 342)

A similar conclusion is reached by British psychologist Guy Claxton, though from the perspective of spiritual practice (see Chapter 18). He argues that much of the trouble in our lives is caused by the false idea of self, and he explores some of the bizarre things we end up thinking when we try to defend the theory that our decisions cause our actions. 'I meant to keep my cool but I just couldn't. . . I'd decided on an early night but somehow here we are in Piccadilly Circus at four a.m. with silly hats and a bottle of wine' (1986a, p. 59). Then if all else fails we can even reinterpret our failure as a success. "I changed my mind", we say, temporarily withdrawing our identification from the "mind" that has been "made up", and aligning ourselves instead with some higher decision-maker and controller who can "choose" to override this mind' (pp. 59–60). But there is no self who really has this control, says Claxton. Rather like Haggard (see Concept 9.1), he concludes that it makes better sense to see the relationship between thought and action as a hit-and-miss attempt at *prediction* rather than control.

The idea that we predict rather than controlling what we do links to a suggestion by Austrian-American psychologist George Mandler (2007) that instead of trying to distinguish between voluntary and involuntary, we should think of a continuum from expected to unexpected. Things seem voluntary merely because they don't surprise us.

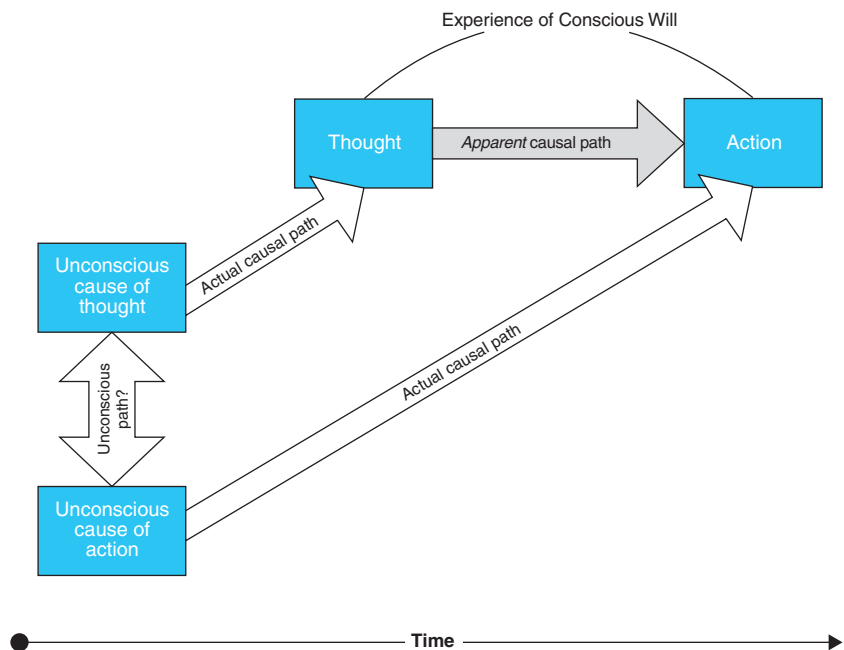


FIGURE 9.8 • According to Wegner, the experience of conscious will arises when a person infers a causal path from thought to action. Both thought and action are caused by unconscious mental events, which may also be linked to each other. The path from thought to action is apparent, not real (after Wegner, 2002, p. 68).

So, is free will always an illusion? Whether free will is what it seems or not, we may draw one firm conclusion. The fact that we may *feel* as though we have free will is not convincing evidence either way.

THE CONSEQUENCES OF BELIEF

A common argument against taking a deterministic view of the universe is that we need belief in the possibility of exercising conscious will to stop us behaving immorally. But does believing in free will really make a difference to how we act, or is the belief that it must make a difference just one more aspect of our illusions about consciousness?

Levels of belief in free will are high in the few surveys that have been done, with scales being developed to measure constructs such as scientific determinism, fatalistic determinism, and perceptions of the world's unpredictability (Rakos et al., 2008; Paulhus and Carey, 2011), but reactions can vary with context. For example, if you ask people whether someone can be free and morally responsible in a deterministic world, they usually say no. But if you ask people whether John, who murdered his wife and children so he could be with his lover, can be free and morally responsible in a deterministic world, they usually say yes. This effect has been repeated with similar results across a number of different deterministic scenarios, and in different languages and cultures (Sarkissian et al., 2010).

This difference has been attributed to emotional reactions to John and his behaviours, which are absent in the abstract case. But a recent meta-analysis of thirty studies found that the size of such emotional reactions is not large enough to explain the effect (Feltz and Cova, 2014). Another possibility is that the mental states of the protagonists are bypassed in the abstract case but may be explicitly given as a cause of action (John wanted to be with his lover) in the concrete case, in a way that mirrors how we think about our own motivations and behaviours. This means that small details of phrasing can make crucial differences in how people interpret the statements. If they read a sentence as implying people cannot act on the basis of their mental states, they give what appear to be incompatibilist answers. But this arguably has nothing to do with incompatibilism, and everything to do with simply not believing that free will is possible if mental states have no impact on action.

Despite these difficulties, we can at least conclude that belief in free will is widespread, but does this belief have consequences for behaviour? It might seem that we could find out by comparing the actions of those who do and do not believe in free will. But this will provide only correlations and not evidence for causality. For example, people who tend to cruel or criminal behaviour might be inclined to reject free will in order to claim that 'my genes made me do it' or 'I couldn't help lying' to avoid the consequences of their actions. Religious believers may behave better because they believe in hell. What we need is experiments in which belief is manipulated.

Many such experiments have been done, mostly priming participants by asking them to read statements provoking either determinist or free-will beliefs. Some have used sections from Crick's *The Astonishing Hypothesis* (1994) (see our Chapter 7). Others give one group of participants such statements as 'Science has demonstrated that free will is an illusion', or 'Like everything else in

the universe, all human actions follow from prior events and ultimately can be understood in terms of the movement of molecules; while another group reads something like, 'I have feelings of regret when I make bad decisions because I know that ultimately I am responsible for my actions', or texts unrelated to free will.

In one such study, those who read about determinism were more likely to cheat on a maths test (Vohs and Schooler, 2008). That these effects were really due to the manipulation was supported by the finding that professed belief in free will was reduced after the reading, and this was correlated with the cheating behaviour. A second study also used pro-free will statements such as 'I am able to override the genetic and environmental factors that sometimes influence my behavior' and 'Avoiding temptation requires that I exert my free will'. People who read these were less likely to overpay themselves for performance on a cognitive task than those who read the pro-determinism statements.

Another set of experiments tested whether inducing deterministic beliefs would induce a 'don't bother' attitude, undermine a sense of responsibility, reduce helping behaviour, and increase aggression (Baumeister et al., 2009). Those who read pro-free will statements did report more willingness to help others and less aggression, but the results suggested the effects were not due to either increased energy to act or an increased sense of responsibility.

A closer look at the mechanisms involved suggests that when people are induced to disbelieve in free will, low-level sensorimotor effects can take place even if people's explicit ratings of sense of agency are unchanged. These include changes in intentional binding (perceptions of how close in time an action seems to its effects; see Concept 9.1), post-error slowing, action-cancellation, and motor preparation for action (Lynn et al., 2014). So, beliefs might intervene at the sensorimotor level and then have a cascade of further effects: on the level of intentional effort exerted, and in turn on our pre-reflective sense of agency and responsibility, regardless of how we report on it.

What is going on here? All this empirical work makes clear that our sense of agency is not a unitary thing, but consists of many different components, of which belief is just one. We may doubt whether these brief experimental manipulations really change people's beliefs in a meaningful way. Thinking about free will and determinism for half an hour is a far cry from the life-long training that some people undertake when once they come to the conclusion that free will is illusory (Blackmore, 2013).

If we accept the findings, however, can we conclude that believing in free will is essential to maintain moral behaviour or even that encouraging people to give up such belief (perhaps by acquainting them with the evidence in a chapter like this one) is bound to lead to unethical behaviour and the breakdown of civil society? This is an argument with a long history. The sixteenth-century Catholic theologian Erasmus wrote that an educated elite might be able to cope with the dangerous idea that there is no free will but the general public was too weak or ignorant to handle such knowledge (1524/1999, pp. 11–12).

My message to you is this: pretend that you have free will. It's essential that you behave as if your decisions matter, even though you know that they don't. The reality isn't important: what's important is your belief, and believing the lie is the only way to avoid a waking coma. Civilization now depends on self-deception. Perhaps it always has.

(Ted Chiang, 'What's Expected of Us', *Nature*, 436, 150 [2005])

Should we then 'protect' people from such dangerous knowledge? Concerned that 'advocating a deterministic worldview could undermine moral behavior', Kathleen Vohs and Jonathan Schooler (2008, p. 54) suggest that 'identifying approaches for insulating the public against this danger becomes imperative'. Wegner himself seemed to share such fears, saying that doubting free will makes everyone uncomfortable and that 'sometimes how things seem is more important than what they are' (2002, pp. 336, 341).

If free will, as commonly conceived, is really illusory, this attitude amounts to a conflict between truth and expediency, with some wanting to keep the truth from people through fear of the consequences. Dennett argues that anyone who gives up free will 'is essentially disabled as a chooser' and that 'the experience, however brief, is grim. And its implications if we take it seriously are almost too grim to contemplate' (1984/2015, p. 184). In *Freedom Evolves* (2003) he propounds his strong compatibilist view, and in a review of Sam Harris's *Free Will* (2012) he lists the dangers he sees in giving up belief in free will. 'If nobody is responsible, not really, then not only should the prisons be emptied, but no contract is valid, mortgages should be abolished, and we can never hold anybody to account for anything they do' (Dennett, 2014). Responding with a blistering attack, Sam Harris (2014) concludes, 'I have not argued for my position primarily out of concern for the consequences of accepting it. And I believe you have'.

Are the consequences really as disastrous as Dennett claims? No, not necessarily. People could still be sent to prison either as a deterrent or in extreme cases to keep everyone else safe. People (whole human beings) can still be held to account for their actions and sign mortgage applications without having to believe they were truly free to do so. Unfree choices (which means all choices if you give up believing in free will) still have consequences and legal implications. And there may be positive consequences such as encouraging compassion for the poor and the mentally ill and discouraging retribution in legal contexts (Greene and Cohen, 2004; Miles, 2013; Shariff et al., 2014).

Independent researcher James Miles argues that the incoherence of many philosophers' and psychologists' positions on free will is actually helping keep our world unequal and unjust. For him, Wegner's statement that the illusion of free will makes us 'who we are' (Wegner, 2002, p. 238) is darkly ironic, because 'the myth of free will does not just excuse indifference to poverty, it creates and maintains much of that poverty in the first place' (2013, p. 216). In his criticism of 'everything that has ever been written by academic philosophers, scientists, and theologians in defence of the notion of free will' (p. 206, see Figure 9.9), Miles includes the category of 'free will illusionism': understanding that free will does not exist but openly misleading

'the responsibility of free will is necessary for belief in a just world'

(Carey, 2009, pp. 8, 20)

'the myth of free will does not just excuse indifference to poverty, it creates and maintains much of that poverty in the first place'

(Miles, 2013, p. 216)

	Can free will and determinism co-exist?	Is determinism true (at the human level)?	Do we have free will?
1. Illusionism	No	Yes	No, but don't tell anyone
2. Compatibilism	'Yes'	Yes	'Yes' (but not free choice)
3. Libertarianism	No	Yes	Yes, but we have no proof

FIGURE 9.9 • Summary of arguments for free will (illusionism, compatibilism, libertarianism) (from Miles, 2013, p. 206).

'when I go to the restaurant and I look at the menu, I might decide "Well, I'll have the spaghetti", but I'm not forced to have the spaghetti; [...] I could have done something else'

(Searle, in Blackmore, 2005, pp. 204–205)

'I do the "as if". And I think almost everybody who's happy and healthy tends to do that.'

(Wegner, in Blackmore, 2005, p. 257)

'the dreaded mayhem does not happen'

(Claxton, 1986a, p. 69)

the public over its non-existence. Many of the researchers discussed here fall into this category. A large part of the problem, he says, is a confusion between determinism and fatalism. Fatalism is the belief that because everything is determined, it is pointless to act. But a determinist, Miles reminds us, will make as many decisions

as a believer in free will; the only difference is that the determinist will recognise her decisions as fully determined. In a restaurant,

The determinist will still select the fish over the wood pigeon, he or she just will not cast the runes seeking instruction, offer up a quick prayer for guidance, or invoke this as proof of either God or free will.

(Miles, 2013, pp. 214–215)

Some take up the challenge of embracing determinism without fatalism when following a spiritual path: the surrendering of will forms part of the mystical traditions of both Christianity and Islam, and Buddhist teachings include the concept of *anatta*, or no-self, which rejects the idea of any persisting entity that acts, and encourages a way of non-action or not-doing (Chapter 18). In his classic book *The Way of Zen*, Alan Watts describes the consequences.

We just decide without having the faintest understanding of how we do it. In fact it is neither voluntary nor involuntary. [...] a decision – the freest of my actions – just happens like hiccups inside me or like a bird singing outside me.

(Watts, 1957, p. 141)

This echoes James's simple 'we *have* got up'. To live this way, it must be 'clear beyond any shadow of doubt that it is actually impossible to do anything else', says Watt (p. 161). This is 'unmotivated non-volitional functioning'. It is how things are because really there is no entity to act, no entity to be either bound or free (Wei Wu Wei, 2004).

Is such complete giving up of free will possible for ordinary mortals? Searle claims not. 'We cannot get rid of the conviction that we are free even if we become philosophically convinced that the conviction is wrong' (2004, p. 219). Interviewing philosophers, psychologists, and neuroscientists about their beliefs, Sue (Blackmore, 2005) found that even those who did not believe in free will often claimed that to live healthily and happily they had to separate their intellectual understanding from the rest of their life, and live 'as if' they did believe.

So these fears run deep, but are they valid? The answer from those who have tried is that the long path to giving up free will leads not to immorality but to kindness, compassion, and personal happiness.

The thing that doesn't happen, but of which people are quite reasonably scared, is that I get worse. A common elaboration of the belief that control is real [...] is that I can, and must control 'myself', and that unless I do, base urges will spill out and I will run amok.

(Claxton, 1986a, p. 69)

Luckily, says Claxton, this is untrue because *I* never was split into controller and controlled, even if the struggle and self-recrimination were real enough. 'So the dreaded mayhem does not happen. I do not take up wholesale rape and pillage and knocking down old ladies just for fun' (p. 69). Instead guilt, shame, embarrassment, self-doubt, fear of failure, and much anxiety fall away, and contrary to expectation I become a better neighbour.

Harris has a similar reaction.

Speaking from personal experience, I think that losing the sense of free will has only improved my ethics – by increasing my feelings of compassion and forgiveness, and diminishing my sense of entitlement to the fruits of my own good luck.

(2012, p. 45)

If they are right, there is no need to protect anyone from anything, and we can welcome the evidence whether it suggests that free will is a genuine force or an illusion.

If you suspect that free will is an illusion, what can you do about it? You can ignore the feeling and hope it will go away. You can act 'as if' you had free will. Or you can stop believing in it. If you choose the third option, you can be sure that everything about your conscious experience will change.




FIGURE 9.10 • Remember that an illusion is not something that does not exist but something that is not what it seems. In this visual illusion the upper monster seems far bigger and more frightening than the lower monster. In fact they are identical. Is consciousness what it seems to be? Is free will?

READING

Dennett, D. C. (1991). Time and experience [excerpt]. In D. C. Dennett, *Consciousness explained* (pp. 153–162). London: Little, Brown.

Dennett's criticism of Libet's notion of backwards referral as a Stalinesque method for establishing the time of consciousness (see Chapter 6).

Haggard, P. (2008). Human volition: towards a neuroscience of will. *Nature Reviews Neuroscience*, 9, 934–946.

Makes the empirical and theoretical case for treating voluntary action, including conscious experience, as something the brain does.

Libet, B. (1985). Unconscious cerebral initiative and the role of conscious will in voluntary action. *Behavioral and Brain Sciences*, 8, 529–539. Commentaries following Libet's article: *BBS*, 8, 539–566, and *BBS*, 10, 318–321 (especially Breitmeyer, Latta, Nelson).


A classic which rewards careful study – both of Libet's original methods and conclusions, and of the many interpretations of these by others.

Miles, J. B. (2013). 'Irresponsible and a disservice': The integrity of social psychology turns on the free will dilemma. *British Journal of Psychology*, 52, 205–218.

Argues that the scientific study of free will is generally biased towards assuming that believing in it is good and not believing is dangerous, when in fact the opposite is true.

Wegner, D. M. (2003). The mind's best trick: How we experience conscious will. *Trends in Cognitive Sciences*, 7, 65–69.


Why we feel we consciously will our own actions, whether or not we really do.



Evolution

SECTION

Four



Evolution and animal minds

CHAPTER

TEN

Humans are animals, so any question about human consciousness is also a question about the animal kingdom as a whole. If there is such a thing as human consciousness, when, why, and how did it come into being, and was the same true for any other animals? It is easy to assume a gradient of consciousness with humans at the top and simple organisms off the bottom of the scale, but whether this is right or wrong has serious implications for how we treat other creatures, as well as for how we understand our own consciousness. This chapter will introduce the basics of evolutionary theory as a foundation on which to ask about the evolution of consciousness (the what-it's-like) in different species, asking where their forms of consciousness may be similar to ours, and where they differ, and what that may mean.

MINDLESS DESIGN

Suppose you are walking along a deserted sandy beach when you come across a magnificent pile of sand. At each corner is a square tower, decorated with rows of shells, and all round is a moat with a flat stone for a bridge, neatly attached to threads of seaweed for pulling it up. What will you conclude? It's a sandcastle of course. And somebody must have built it.

When we see obvious signs of design, we readily infer a designer. This, in essence, is the 'argument from design', made famous in 1802 by the Reverend William Paley. He supposed that, crossing a heath, he found either a stone or a watch. For the stone, he could conclude that it had always been there, but for the watch he *must* conclude that it had a maker. All the parts are ingeniously linked to serve the

*'There cannot be design
without a designer'*

(Paley, 1802, p. 3)

• SECTION FOUR : EVOLUTION

purpose for which it was constructed – telling the time. If any pieces were missing or in the wrong place or material, the watch would not work. He could not see how these many complex pieces could have come together by accident, nor by the effects of natural forces such as wind or rain, so he concluded

that the watch must have had a maker: that there must have existed, at some time, and at some place or other, an artificer or artificers who formed it for the purpose which we find it actually to answer: who comprehended its construction, and designed its use.

(Paley, 1802, p. 3)

He thought it self-evident that ‘There cannot be design without a designer; contrivance, without a contriver; order, without choice’. The arrangement of the functioning parts must ‘imply the presence of intelligence and mind’ (Paley, p. 13).

So it is, he said, with the wonders of nature: the intricate design of the eye for seeing, the ways in which animals attract their mates, the design of valves to aid the circulation of the blood – all these show complex design for a purpose and hence they must have had a designer. In this way, the argument from design becomes evidence for the existence of God.

Paley was wrong. We now understand, as he could not have, that design does not need a designer. As Oxford biologist Richard Dawkins puts it, ‘Paley’s argument is made with passionate sincerity and is informed by the best biological scholarship of his day, but it is wrong, gloriously and utterly wrong’ (Dawkins, 1986, p. 5). There are not just two possibilities – accident, or design by a conscious intelligent designer. There is a third that no one could have known about in Paley’s day. Design for function can appear without a designer, and Darwin’s theory of evolution by natural selection showed how.



FIGURE 10.1 • These are some of the finches that Darwin collected in the Galapagos Islands in 1835. Each species has a different shape of beak – essentially a tool designed for a specific job, from picking tiny seeds out of crevices to crushing nuts or shells. At the time it seemed obvious that God must have designed each one. As Darwin put it in his memoir *The Voyage of the Beagle* (1839/1909, p. 402), ‘one might really fancy that . . . one species had been taken and modified for different ends’. But in 1859 Darwin explained how beaks, finches, and the entire natural world could have been designed without a designer – by natural selection.

‘Evolution’ means gradual change, and the idea that living things might, in this general sense, evolve was already current in Darwin’s time. His own grandfather, Erasmus Darwin, had questioned the prevailing assumption that species were fixed by God. And Sir Charles Lyell’s theory that geological forces could carve landscapes, shape rivers, and throw up mountains already threatened the idea that God had designed the earth just as we find it today. The fossil record suggested gradual change in living things, and this demanded explanation. What was missing was any mechanism to explain how evolution worked. This is what Darwin provided, in his 1859 book *The Origin of Species*. Its full title is *On the Origin of Species by Means of Natural Selection, or the Preservation of Favoured Races in the Struggle for Life*.

His idea was this. If, over a long period of time, creatures vary (as he showed they did), and if there is sometimes a severe struggle for life (which could not be disputed – he had read Malthus’s *Essay on Population*), then occasionally some variation in structure or habits must occur that is advantageous to a creature. When this happens, individuals with that characteristic have the best chance of being

preserved in the struggle for life, and they will produce offspring similarly characterised. This ‘principle of preservation, or the survival of the fittest’, he called ‘Natural Selection’. It leads to the improvement of each creature in relation to its conditions of life.

In more modern language, we might put it this way. If many slightly different creatures have to compete for food, water, or other resources, and many die, and if the survivors pass on whatever helped them survive, then their offspring must be better adapted to that environment than their parents were. With long repetition of selection over billions of years, extraordinary adaptations can gradually appear, including fur, legs, wings, and eyes.

Paley was especially concerned with eyes, because of their intricate and delicate design, but the principle is just the same for eyes as for anything else. In a population of creatures with single photosensitive cells, those with more cells might have an advantage; in a population with eye pits, those with deeper pits might do better; and so on until eyes with corneas, lenses, and foveas are forced into existence. It is now thought that eyes have evolved independently more than forty times on planet Earth. Natural selection is not the only force in evolution, but together with mutation, genetic drift, gene flow, sexual selection, and layers of self-organisation from the molecular level upwards, it explains how design appears naturally without a plan or a designer.

‘Nothing in biology makes sense except in the light of evolution’, proclaimed biologist Theodosius Dobzhansky (1973). Natural selection is ‘one of the most powerful ideas in all science’ (Mark Ridley, 1996, p. 3), ‘the single best idea anyone has ever had’ (Dennett, 1995b, p. 4). ‘Darwin’s Dangerous Idea’ is like a universal acid that eats through everything in its path, revolutionising our world view as it goes (1995b, Ch. 3). This ‘dangerous idea’ has become the foundation for all the biological sciences.

The process that Darwin described as ‘descent with modification’ can be thought of as a three-step algorithm: if you have variation, heredity, and selection, then you *must* get evolution. It is ‘a scheme for creating Design out of Chaos without the aid of Mind’ (Dennett, 1995b, p. 50; see also pp. 48–52, 61–89, 324–330, and 521). American psychologist Donald Campbell (1960) described it as ‘blind variation and selective retention’. Since clever designs thrive because their competitors don’t, we could also think of it as ‘design by death’.

This scheme is an inevitable process that requires no designer and no plan. It needs no foresight and no intentions. It need not happen for any purpose or towards any end. It could all be done by a ‘Blind Watchmaker’ (Dawkins, 1986). Paley’s eyes and ears, valves and mating calls were designed all right, but no designer was required.

DIRECTED EVOLUTION

Despite Darwin’s insight, the idea that evolution still requires a guiding hand seems endlessly appealing and has often reappeared. Jean-Baptiste Lamarck (1744–1829) agreed with Darwin that species might gradually change into other species, but he proposed first an individual force (an animal’s drive to adapt to its conditions) that produced progress in one direction, and second the inheritance of acquired characteristics (this is now referred to as Lamarckism even though

If you have **Variation**
Selection and
Heredity
You must get **Evolution**
Or ‘Design out of Chaos
without the aid of Mind’

FIGURE 10.2 • The evolutionary algorithm (Dennett, 1995b, p. 50)

‘Nothing in biology makes sense except in the light of evolution’

(Dobzhansky, 1973)

Natural selection is ‘a scheme for creating Design out of Chaos without the aid of Mind’

(Dennett, 1995b, p. 50)




FIGURE 10.3 • Victorians were scandalised by Darwin's suggestion that civilised human beings might be related to the apes. He was mocked and lampooned, as in this cartoon from the *London Sketch Book* of 1874.

Lamarck was not the first to suggest it, and Darwin wrote about similar processes). Lamarck believed that if an animal used a particular faculty to change itself, the effect would be passed on to its offspring. So a giraffe that spent its life stretching to the highest branches would have calves with slightly longer necks; a blacksmith who worked hard and developed huge muscles would pass on the effects to his children.

These two theories provide very different visions of evolution and its future. On Lamarck's scheme, evolution is directional and progressive, with species inevitably improving over time. On Darwin's scheme there is no guarantee of progress and no inbuilt direction. The process produces a vast tree or straggly bush of species and subspecies, with branches appearing all over the place, change always starting from whatever is available, and species going extinct when conditions dictate. Darwin's scheme has no special place for humans, who are just one chance product of a long and complex process, rather than its inevitable outcome or highest creation.


Not surprisingly, Lamarck's vision proved more acceptable than Darwin's and is still popular today. Darwin's faced massive resistance from religion, and was met with ridicule and contempt. At a famous debate in Oxford in 1860, the Bishop of Oxford, Samuel Wilberforce, asked Thomas Henry Huxley, Darwin's main protagonist, whether he was descended from the apes on his grandmother's side or his grandfather's side, to great popular amusement. Even today there is religious opposition to Darwinism in some countries, including the United States of America, where the idea of directed evolution underlies both creationism and its successor 'intelligent design', with God as the supreme director who creates human beings 'in His image'.

The 'Great Chain of Being' is another alluring idea, with simple organisms at one end and conscious, intelligent human beings at the other. So is the image of an evolutionary ladder with humans striving to climb from lowly creatures at the bottom to angels at the top. Such schemes seem

to justify our struggles and imply that progress is directed by our efforts. Lamarck's views have often been interpreted as meaning that those efforts involve consciously willed striving. This is not what Lamarck said, even though he gave much thought to how physiological processing gives rise to 'inner feeling', or conscious experience. But since then, many theories have given a more explicitly central role to consciousness. For example, the Jesuit priest Pierre Teilhard de Chardin (1959) proposed that all life is striving towards higher consciousness, or the 'Omega Point'. For the Indian visionary Sri Aurobindo, life is evolving into the 'life divine', and biologist Julian Huxley believed that evolution has become truly purposeful and 'is pulled on consciously from in front as well as being impelled blindly from behind' (in Pickering and Skinner, 1990, p. 83).

Some modern 'spiritual' theories also invoke conscious direction, such as Ken Wilber's 'integral theory of consciousness'. This theory is explicitly based on the great chain of being, and on the idea of inevitable progress from insentient matter to superconsciousness or transcendence (Wilber, 1997).

Maybe the evolutionary sequence really is from matter to body to mind to soul to spirit, each [. . .] with a greater depth and greater consciousness



and wider embrace. And in the highest reaches of evolution [. . .] a Kosmic consciousness that is Spirit awakened to its own true nature.

(Wilber, 2001, p. 62)

Wilber explicitly rejects the evidence accounting for the evolution of wings and eyes, arguing, as do creationists, that wings could not have evolved naturally because half a wing or half an eye would be of no use.

Three other popular examples are the motivational theory of futurist Barbara Marx Hubbard (1997), who urges us all to realise the potential of our higher consciousness and take control of our own future in 'conscious evolution'; the quantum theory of physicist Amit Goswami (2008), who argues that consciousness, rather than matter or energy, is the primary force in the universe; and the ideas of Deepak Chopra (Chopra and Tanzi, 2012), a practitioner of integrative medicine who believes that Darwin was wrong because a supernatural consciousness directs evolution and allows humans to escape the forces of natural selection that other animals are bound by. All three follow Teilhard de Chardin in believing that evolution is driven from above by a cosmic or spiritual field. The idea that consciousness drives evolution forward seems to have timeless appeal, but so far there is no biological evidence providing any reason to believe in it.




FIGURE 10.4 • In the popular idea of a 'great chain of being', evolution proceeds through a line of ever-improving creatures to culminate in the most perfect and intelligent of them all – 'man'. The reality is more like a branching tree, or a great bush, in which humans are on one twig of the primate branch. On this view, all the creatures alive today are adapted to their niche, and none is necessarily 'more evolved' or 'higher' than the rest.

The reason for rejecting Lamarckism (at least the popular version of it) was first made clear by August Weismann (1833–1914), who distinguished, in sexual species, between the germ line (the sex cells that are passed from generation to generation) and the soma (the body which dies). What happens to a body affects its *chances* of passing on its sex cells, but not those cells themselves. Nowadays we would say that genetic information (the genotype) is used to construct the body (the phenotype), and that changes to the phenotype cannot affect the genotype. So, for example, if you spend your life dieting you may make yourself more or less attractive, or even infertile, but you will not pass on genes for slimmer children.

We now know, however, that the food you eat, and other lifestyle choices, can and do have effects on future generations through epigenesis. For example, a famous study found that children of pregnant women who lived through the Dutch famine during the Second World War were more susceptible to obesity, diabetes, heart disease, and other health problems in later life and that these effects may even have been passed on to the next generation (Veenendaal et al.,

2013). Epigenetic effects do not actually change any genes but involve heritable changes in the way genes are expressed, including switching them on and off.

Yet the basic distinction between germ line and soma remains, and there are good reasons for evolution to work this way. Most things that happen to phenotypes are harmful, such as failures in development, damage of various kinds, and ageing. If all these changes were passed on, useful developments in design would be lost. Also, most phenotypes are just not very successful, so it makes sense to go 'back to the drawing board' in each generation (Dawkins, 1989). Another way of putting it is that schemes which copy the instructions for making a product (such as making organisms from instructions in the DNA, or building cars on a production line) are better than schemes that copy the product itself because of the inevitable errors introduced by imperfect copying (Blackmore, 1999).

By the early twentieth century, Darwinism was in the doldrums, but change came in the 1930s with the discovery of the basis of genetics and its integration with natural selection in 'the modern synthesis'. The resulting neo-Darwinism explained why no directing force was needed; natural selection working on variation created by the recombination and mutation of genes was sufficient. Clearly other processes such as genetic drift, gene flow, random events, epigenetic inheritance, and self-organising principles play a role in evolution, and there have been fierce arguments over their relative contributions (R. Dawkins, 1986; Dennett, 1995b; Gould and Lewontin, 1979; Jablonka, Lamb, and Zeligowski, 2005; Johnson and Lam, 2010). Even so, there is no hint of a guiding force in evolution and no evidence to suggest that mind or consciousness plays that role.

SELFISH REPLICATORS

Who or what is evolution for? Who or what is the ultimate beneficiary of eyes, wings, brains, and digestive systems? Darwinism is frequently misunderstood as a mechanism that creates adaptations 'for the good of the species'. A simple example will show why it is not.

Imagine a population of rats successfully living off human rubbish in a huge modern city – let's say London. Outside every shop and restaurant are plenty of dustbins that contain plenty of nice rat food. Every night when the workers leave, there is a chance that the dustbins will not be properly sealed, or food will be left on the ground. As long as the rats wait quietly until the humans have left, they will have it all to themselves. The best strategy 'for the good of the species' is for every rat to wait, but will they? Of course not. If just one rat has genes that incline it to jump in first, causing the dustbin lid to clatter to the ground and the humans to come running to close it, that rat will still be better off than the rest, running off with some nice rotting meat or a soggy sandwich. That rat will get fatter, take more food home, and produce more offspring, who will also tend to inherit the 'jump first' tendency. The patient rats lose out. Note that this general point is not a recipe for unadulterated selfishness. There are many reasons why cooperative and altruistic behaviours can thrive alongside selfish ones (Matt Ridley, 1996; Fletcher and Doebeli, 2009; Nowak and Highfield, 2011). It is, however, an argument against evolution proceeding 'for the good of the species'. We must not, therefore, fall into the error of thinking that consciousness could have evolved because it was good for our species, or indeed for any other species.

So is the individual the ultimate beneficiary? Not necessarily. In his classic 1966 book *Adaptation and Natural Selection*, the American biologist George Williams argued that we should recognise adaptations at the level necessitated by the facts and no higher. But which level is that? Multilevel selection theory entails selection operating at many levels, including group selection in which groups of animals, or tribes, or cultures compete with each other for survival. Arguably, cultural evolution makes this more likely (Boyd and Richerson, 2009). If, for example, one human group has a taboo on eating pork in a region where pigs carry fatal diseases, that group would have an advantage. In this case, selection would operate both within the groups and between them. But group selection is a highly contentious topic. Powerful advocates include David Sloan Wilson (Wilson and Sober, 1994) and E. O. Wilson (D. S. Wilson and E. O. Wilson, 2008), while Steven Pinker (2016, p. 878) says it is guaranteed to confuse because it is too often used to make 'loose allusions to the importance of groups in human evolution'.

Against group selection is what is known as 'selfish gene theory' after Dawkins's 1976 book *The Selfish Gene*. On this view, the ultimate beneficiary of natural selection is neither the species, nor the group, nor even the individual, but the hereditary information: the gene. If this seems odd, think about our London rats. They have genes for numerous physical and behavioural traits, and natural selection can work on them all. Although it is the individual rats who live or die, the net result is changes in the frequency of different genes in the gene pool. Another way of putting it is to say that the gene is the 'replicator': it is the information that is copied, either accurately and frequently, or not. This explains how genes can be 'selfish'. They are *not* selfish in the sense of having their own desires or intentions (they couldn't; they are just information coded on strands of DNA); they are *not* selfish in the sense that they produce selfish behaviour in their carriers (they produce altruistic behaviours too); but they *are* selfish in the sense that they will get copied if they can – regardless of their effect on other genes, on their own organisms, or on the species as a whole. From this perspective, human beings (like all other animals) are the 'lumbering robots' that have been designed by natural selection to carry the genes around and protect them (Dawkins, 1976).

There is a danger of seeing every trait as necessarily adaptive (a tendency derided as 'panadaptationism' by palaeontologist Stephen Jay Gould [Gould and Lewontin, 1979]). In fact, many features of organisms are not adaptations, or are far from optimal if they are. Some are strongly influenced by physical constraints and random forces, and none is optimally designed because evolution always has to start from whatever is available and work from there. Some useless traits survive



PROFILE 10.1

Richard Dawkins (b. 1941)



Born in Nairobi, Dawkins came to England with his family in 1949. He studied at the University of

Oxford, where he subsequently became Lecturer in Zoology, Fellow of New College, and then Charles Simonyi Professor of the Public Understanding of Science until his retirement in 2008. His first book, *The Selfish Gene* (1976), established what came to be called 'selfish gene theory', and was a bestseller for many decades. As a protagonist in the 'Darwin Wars', he battled against Stephen Jay Gould over the importance of natural selection and adaptation in evolution (Brown, 1999; Sterelny, 2001). His book *The God Delusion* (2007) inspired 'the new atheism', a movement against religious dogma and indoctrination whose main proponents, including Dawkins, were dubbed 'The four horsemen'. He describes human beings as mere 'survival machines' – the 'lumbering robots' designed to carry our genes around. In promoting 'Universal Darwinism', he invented the concept of the meme as a cultural replicator, and refers to religions as viruses of the mind. As for consciousness, he thinks it is 'the most profound mystery facing modern biology'.

'what is a single selfish gene trying to do? It is trying to get more numerous in the gene pool'

(R. Dawkins, 1989, p. 88)

because they are by-products of other traits that have been selected for. Others survive because they were once adaptive and there has not been sufficient selection pressure to weed them out. All these may be possibilities when we ask *why* consciousness has evolved – a question to which we return in Chapter 11.

ANIMAL MINDS

What is it like to be one of our successful London rats? Or a snake? Or a goldfish in a tank? Or a butterfly? We cannot think about the evolution of human consciousness without also asking about other animals. The human lineage is thought to have split from that of chimpanzees between 5 and 7 million years ago, from gorillas between 8 and 10 million years ago, and from orangutans between 12 and 16 million years ago. Human DNA is approximately 94.8% identical to that of chimpanzees, so are the other great apes conscious too? Are monkeys? Are squirrels? Is there something it's like to be a grey squirrel burying some hazelnuts for the winter? If not, what change occurred to make humans conscious and leave other species – even other primates – 'in the dark'? On the other hand, if gorillas are conscious, is a single-celled organism? If so, why it and not the complex lattice structure of diamond? Or does asking this kind of question reveal, above all, how easily we tie ourselves in knots when thinking about consciousness?

It is easy to imagine a ladder in which humans, at the top, have the highest levels of consciousness – or are the only ones who are conscious – while further down consciousness is different, more impoverished, or absent altogether. But is there any evidence for this kind of model? In the rest of this chapter, we consider a range of methods we can use to investigate whether other animals are conscious – and if so, how their conscious experiences compare to ours. We will try to resist the seduction of the linear scale with us at the top, though inevitably much of the research asks which of our human skills other animals are capable of and what this says about their capacity for consciousness.

I had of course long been used to a halter and a headstall, and to be led about in the fields and lanes quietly, but now I was to have a bit and bridle; my master gave me some oats as usual, and after a good deal of coaxing he got the bit into my mouth, and the bridle fixed, but it was a nasty thing! Those who have never had a bit in their mouths cannot think how bad it feels; a great piece of cold hard steel as thick as a man's finger to be pushed into one's mouth, between one's teeth, and over one's tongue, with the ends coming out at the corner of your mouth, and held fast there by straps over your head, under your throat, round your nose, and under your chin; so that no way in the world can you get rid of the nasty hard thing; it is very bad! yes, very bad! at least I thought so; but I knew my mother always wore one when she went out, and all horses did when they were grown up; and so, what with the nice oats, and what with my master's pats, kind words, and gentle ways, I got to wear my bit and bridle.

(Anna Sewall, *Black Beauty: The autobiography of a horse*, 1877)

Let us begin with a thought experiment. What is it like to be an octopus? Can you imagine how it feels to swim swiftly underwater trailing your eight long tentacle arms behind you, using your many suction pads to explore a coral reef? Can you imagine having no skeleton to prevent you squeezing into tiny gaps between rocks, and being able to spray thick dark ink in a big cloud around you to confuse your predators? Maybe you can. But, as Nagel (1974) pointed out in ‘What is it like to be a bat?’ (Chapter 2), you are probably imagining what it would be like for *you* to be the octopus, and that is not the point. The point is what it is like for the octopus – that is, if it is like anything at all for the octopus.




FIGURE 10.5 • *Octopus vulgaris* is a marine cephalopod that uses its arms with two rows of suckers on each to move across and grasp objects. It hunts at dusk, using nerve poison in its saliva to paralyse its prey, and grasping prey with its powerful arms. It is intelligent enough to unscrew jars and raid lobster traps, it can squeeze through small gaps and can change colour to blend in with its surroundings, and uses its light-sensitive skin to detect changes in brightness without using its eyes. Males use the tip of their third right arm to insert sperm into the oviducts of females.

How can we ever know? We cannot ask the octopus to tell us. And even if we could, we might not believe, or understand, what it said. This is essentially the problem of other minds. Just as you can never be sure whether your best friend is really conscious, so you can never know whether your cat, or the birds in your garden, or the ant you just stood on, are (or were) conscious. Humans and other animals show similar expressions of emotion, and similar reactions to pleasure, pain, and fear, as Darwin (1872) long ago showed. From these similarities, we can guess what another animal is trying to do or how it is feeling. Even so, we must avoid assuming that just because it appears to be in pain, or to be feeling guilty, or happy or sad, it really has the feelings we attribute to it. Our impressions could be completely wrong.

There are two extreme positions to consider. One is that only humans are conscious. Descartes believed that because they do not have language, all other animals are unfeeling automata, without souls or consciousness. A modern version is Macphail’s argument that ‘animals are indeed Cartesian machines, and it is the availability of language that confers on us, first, the ability to be self-conscious, and second, the ability to feel’ (Macphail, 1998, p. 233). In his view, there is no convincing evidence for consciousness in other species. They are not just devoid of speech and self-awareness, but devoid of feeling (by which he means sensory experiences), too. Dennett (1991) provides a different reason: that other animals lack the language with which to create the particular kind of fiction that is conscious experience. Similarly, HOT theories deny consciousness to any animal incapable of having higher-order thoughts.

At the other extreme lies the view that all other species are conscious. Panpsychism is the obvious example here: even an amoeba, and beyond that even the inorganic world, has something ‘which is of the same nature with our own consciousness’, although that *something* may be inconceivably simple in comparison (Clifford, 1874/1886, p. 266). Between these extremes lie theories that, for various reasons, attribute different kinds of consciousness to different species (Griffin and Speck, 2004; D. Edelman and Seth, 2009). For example, Baars (2005b) argues


‘It has a body – but one that is protean, all possibility [. . .]. The octopus lives outside the usual body/brain divide’

(Godfrey-Smith, 2017)

‘animals are indeed Cartesian machines’

(Macphail, 1998, p. 233)

1. A dichotomy (one big division):



2. A continuum (seamless transition):



3. A space with many discontinuities:




FIGURE 10.6 • Models of conceptual spaces. It is often assumed that the only alternative to a dichotomy (conscious/nonconscious) is a continuum of cases with only differences of degree. There is a third alternative (Sloman and Chrisley, 2003, p. 15).

'affective, interoceptive, and exteroceptive consciousness all existed in the first vertebrates of the Cambrian explosion'

(Feinberg and Mallatt, 2016, p. xvii)

that consciousness is a fundamental biological adaptation and that the known anatomical and physiological bases of consciousness are phylogenetically ancient, going back at least to the early mammals. We readily attribute consciousness to other people on the basis of behavioural and brain evidence, so we should not deny it to other mammals. Psychiatrist Todd Feinberg and biologist Jon Mallatt (2016) go even further back, to the time of the Cambrian explosion more than five hundred million years ago. Unconscious reflexes gradually evolved into brains with ever increasing levels of consciousness leading eventually to unified inner worlds of subjective experience. So in their view, every fish, reptile, amphibian, and insect is conscious, and possibly cephalopods like our octopus, too.

Others distinguish between primary consciousness, or sensory consciousness – the ability to integrate perceptual and motor events with memory to create awareness of the present world around you – and secondary or higher-order consciousness, which involves consciousness of being conscious, and the ability to connect the present to the past and future (G. Edelman, 2003). Distinctions like these may mean that consciousness is basically a binary, on/off phenomenon, or they may allow for certain animals to be partially or incompletely conscious (Allen and Trestman, 2016). We should not, in any case, assume that there is just one kind of consciousness. Nor should we assume that, if there are multiple kinds or levels of consciousness, the human kind is the standard by which all the others should be measured.



PRACTICE 10.1

WHAT IS IT LIKE TO BE THAT ANIMAL?

This practice is rather different from usual. As you go about your daily life, look out for other animals and watch what they are doing. They might be pet dogs and cats, farm cows or pigs, or wild birds, squirrels, or rabbits. Look out as well for insects, spiders, worms, and fish. In each case ask yourself, 'What is it like to be this cow?', 'What is it like to be that spider?' Can you imagine it? Is it easier with some animals than with others? What does this difference mean?

DIFFERENT WORLDS

Every species has evolved sensory systems to suit its way of life. This leads to the odd realisation that several different species in the same location may all be inhabiting different worlds. Let's take the example of an ordinary garden pond with fish, frogs, newts, snails, insect larvae, flies, and a human child with a fishing net. We can easily imagine (or think we can) how the pond looks to the child, but the others must experience it in completely different ways. The fish have sense organs for detecting vibrations in the water from which they know what to avoid,

what to seek out, and when to dive for safety. We have nothing comparable to help us imagine it. The insects have compound eyes quite unlike our image-forming eyes, and many of the animals have chemical senses far more sensitive than our feeble senses of smell and taste.

The frog is particularly interesting. Frogs have eyes with lenses and retinas somewhat like ours, sending signals along the optic nerve to the optic tectum in the brain. It is tempting to imagine that a picture of the frog's world is somehow constructed in its brain, but this is not so. The frog's eye tells the frog's brain just what it needs to know and no more. It tells it about stationary and moving edges, changes in overall illumination, and bugs. The 'bug detectors' respond specifically to small moving objects, not to large moving ones or small still ones, and direct the frog's tongue to catch flies. An extraordinary consequence of the way this system works is that a frog can literally starve to death surrounded by freshly killed flies. If the fly does not move, the frog does not see it.

We can learn much from thinking about this frog. We might be inclined to think that the child gazing into the pond really does have a picture of the world in her head – a full, rich, and detailed picture of the scene – and that by comparison the frog's vision is simply stupid. But think again. The discoveries of change blindness and inattentive blindness, and of the different roles of the dorsal and ventral streams in the visual system (Chapter 6), suggest that we may be much more like the frog than we care to admit. Evolution has designed us to detect only selected aspects of the world around us, often only when we need them for action. Just like the frog, we are quite unaware of everything else – yet we feel no gaps.

We may think that the child must be more conscious than the frog, and the frog more conscious than the fly, but why? While many authors make bold assertions about animal consciousness, it is not clear how these can be tested or what they mean. British pharmacologist Susan Greenfield proposes that 'consciousness increases with brain size across the animal kingdom' (2000, p. 180). But if she is right, then sperm whales, African elephants and dusky dolphins are all more conscious than you are, and Great Danes and Labradors are more conscious than Jack Russells and Pekinese. Searle claims that 'Humans and higher animals are obviously conscious, but we do not know how far down the phylogenetic scale consciousness extends' (1997, p. 5). But this is not 'obvious', and there is no single phylogenetic scale, or linear sequence, along which animals can be graded from 'higher' to 'lower'. As we have seen, evolution has produced not a line but a very bushy bush.

'He [the frog] will starve to death surrounded by food if it is not moving'

(Lettvin et al., 1968, p. 1940)

'Humans and higher animals are obviously conscious'

(Searle, 1997, p. 5)

PHYSICAL AND BEHAVIOURAL CRITERIA

Ideally we need to find some clear criteria for consciousness that we can apply to other animals. One way is to look for anatomical or other physical features – not just brain size but aspects of brain organisation and function that we think are indicators of consciousness. We might argue that fish cannot be conscious because human consciousness relies on signal amplification and global integration, and fish lack the neural architecture that makes these possible, in particular the strongly interconnected feedforward and feedback circuitry that allows for neural signals to be both differentiated and integrated (Key, 2016).

'What then do noxious stimuli feel like to a fish? The evidence best supports the idea that they don't feel like anything to a fish'

(Key, 2016, p. 17)

'Consciousness probably evolved first in fishes'

(Balcombe, 2016, p. 85)

'we seek the minimum number of levels a sensory hierarchy can have to produce consciousness'

(Feinberg and Mallatt, 2016, p. 98)

Computational neuroscientist Anil Seth and his colleagues (2005) argue that among the basic brain facts are that consciousness 'involves widespread, relatively fast, low-amplitude interactions in the thalamocortical core of the brain, driven by current tasks and conditions' (p. 119). The lower brainstem is involved in maintaining the state of consciousness, while the thalamocortical complex sustains conscious contents. So, finding these features in the brains of other species should show us that they are conscious. Seth concludes that most mammals share these structures and therefore should be considered conscious.

What about those many creatures that have no cortex and therefore no thalamocortical connections, from brainless molluscs, through tiny-brained worms and insects, to fish and reptiles? Bjorn Merker (2007) argues that all vertebrate brains share a centralised functional design with an upper brainstem system organised for conscious function. In simple brains, this system is involved in action control; in more complex ones, it takes on the task of integrating the massively parallel processing of the higher brain areas into the limited-capacity serial processing required for coherent behaviour. On this view, even simple-brained creatures with no cortex at all can be conscious.

A common theme here is that the brainstem controls states of consciousness and the sleep-waking cycle, while the forebrain sustains complex contents of consciousness. All mammals, and most other animals (including many fish and reptiles, some insects, and even the simple roundworm *C. elegans*), alternate between waking and sleeping states, or at least have strong circadian rhythms of activity and responsiveness. So, in the sense of being awake, they are conscious, but is there something it's like to be them: are they having conscious perceptions, thoughts, feelings? When it comes to conscious 'contents', we face again the difficulties involved in pinning down the NCCs and the problems we encountered with the whole notion of the 'contents of consciousness' (Chapter 4). These problems are even more acute when asking about the NCCs of non-human animals. Here, it is even more difficult to distinguish between prerequisites, substrates, and consequences of conscious experience – and, of course, to determine what experiences count as conscious in the first place (Boly et al., 2013).

If we had a complete theory that specified the neural basis of consciousness, we could use it to determine the status of animals' minds. But we do not. As Seth and colleagues (2005) point out, neural theories of consciousness are new, and the list of criteria may need to change. And until then, we should not just guess which features are needed for consciousness and go looking for them. This is what Feinberg and Mallatt (2016) appear to do when they specify that the 'defining features of consciousness' include non-nested and nested hierarchical functions, isomorphic representations, and mental images, and that sensory hierarchies require four or more levels to be conscious. Seeking these in other species is how they arrived at their conclusion that 'the transition from non-conscious to conscious' happened between 560 and 520 million years ago.

The other main approach is to look at behavioural indicators. For example, a mobile lifestyle (octopuses, not clams; animals, not plants) might drive the need for general-purpose perception, flexible planning, and precisely controlled action, and these might be conducive to developing subjectivity (Klein and Barron, 2016). We might also ask whether organisms capable of particular types of

associative learning that have behavioural as well as functional and structural characteristics are more likely to be conscious (Bronfman, Ginsburg, and Jablonka, 2016).

Or we might try to grade animals by intelligence, but one danger is that we base our idea of intelligence on our own species-specific abilities and fail to appreciate other kinds of intelligence, like those of bees or elephants or octopuses (Adams and Burbeck, 2012; Godfrey-Smith, 2016). Even in more familiar creatures, comparisons are difficult. On some scales, chimpanzees are put near the very top and birds, with their tiny 'bird-brains', much lower down. It is true that chimpanzees can work out how to pile up boxes to reach a suspended banana, but then ravens are just as good as the great apes (and small children) at planning ahead in a domain-flexible way for tool use and bartering (Kabadayi and Osvath, 2017). Is one species more intelligent than the other, or more conscious?

And what about suffering? Does one species suffer more than the other? We have empathy for other people when we see them crying or in distress, which may be reasonable on the assumption that they are similar to us. We may also feel empathy for the dog who squeals when hurt, the tiger pacing up and down in a tiny cage, or the lobster screaming in boiling water. But could any, or all, of them be Cartesian automata that feel nothing? This is not an empty question, since we can build a simple toy dog, wired up so that if you stand on its foot it whines, but few would believe it was capable of suffering. A few switches are not sufficient. But what *is* sufficient for the capacity to suffer (Linzey, 2009)?

Does suffering even require a capacity for consciousness? We may assume it does, but the philosopher John Carruthers (2004), using a higher-order account of consciousness, argues that suffering is possible without phenomenal consciousness; most animals are probably not capable of higher-order thoughts and therefore not conscious, but what makes pain awful is the first-order content. Marian Stamp Dawkins (2008) agrees that the problem of consciousness is separate from the problem of suffering: even though we associate human suffering with the subjective experience of emotions, emotions can also be unconscious.

Does any of this allow us to decide which animals are capable of suffering? The argument over lobsters has been especially fierce. The screaming noise they make when boiled alive upsets people but is probably produced by air being forced out of the shell. Since crustaceans have a simple brain with no cortex,



FIGURE 10.7 • Chimpanzees use sticks and leaves as tools, for example for fishing termites out of holes. There is even evidence that in the past they used stone tools, and that females may be more adept with tools than males. Is this intelligent behaviour a sign of consciousness?



FIGURE 10.8 • NC crows make complex hook tools in the wild from twigs, and use the hooks to extract prey from crevices (J. Troscianko and Rutz, 2015) — shown here from a camera hidden inside a tube (J. Troscianko et al., 2012). Does this make them more intelligent than other crow species, or other birds? Does it make them more conscious? Or does it have more to do with being able to hold a tool and see what you're doing with it, and then naturally developing more solutions that depend on tools?

'The scientific study of animal suffering [...] requires the testing of the untestable'

(M. Dawkins, 2008, p. 1)

'we do not have to solve the problem of consciousness to have a science of animal welfare'

(M. Dawkins, 2008, p. 4)

they cannot have the specific cortical areas associated with pain in humans. So some researchers conclude that they cannot feel any pain at all (an anatomical criterion). Others have shown that when crabs and lobsters are lifted out of water (and so become oxygen-deprived), or are subjected to infection with a parasite, or have a claw twisted off, they release a stress hormone similar to cortisone and corticosterone. This kind of stress response suggests a basis for suffering comparable to that of humans (a physiological criterion). It has also been established that when acid is brushed onto prawns' antennae they quickly rub it off, that they avoid situations where they have been given electric shocks in the past, and that they show protective behaviours such as rubbing and limping when hurt. These behaviours are also reduced when they are given painkillers. So they must be capable of feeling pain (behavioural criteria). How can we weigh up these different indicators (Elwood, Barr, and Paterson, 2009) and make sense of what prawns and lobsters might be feeling, if anything?

Dawkins (2008) suggests two questions we can ask to decide whether an animal is suffering: is the animal healthy, and does it have what it wants? For example, working with broiler chickens, she was surprised to find that although the birds' walking ability was worse in the highest-density farms, space was much less relevant to other health measures like mortality and the state of their legs and feet than environmental factors such as air and litter quality. The chickens also did not seem to try to avoid each other, but seemed to positively like being close to others. She argues that animal welfare matters greatly but in trying to understand it, we should stick to the evidence and not be distracted by anthropomorphism, empathy, or arguments about animal consciousness.

Lurking here is the question, does it really *hurt*? This may seem impossible to answer, but we should not despair. In studying human consciousness, we have made progress by learning about perception, memory, attention, and other relevant abilities. Perhaps we can do the same for animal consciousness.

In what follows, we will survey a range of other behavioural routes to trying to pinpoint consciousness in other animals – asking where their forms of consciousness may be similar to ours, where they may differ, and what that may mean.

SELF-RECOGNITION

You are aware not only of the world around you, but of yourself as an observer. You are self-conscious. It is hard to determine when young children first become self-conscious, but by 5–6 months, infants shown a video of another same-age infant are more captivated by it than by a video of themselves wearing the same clothes. This does not yet mean that they *recognise* themselves; just that they have learned to pick up the invariant features of their own faces and bodies, presumably via exposure to



PROFILE 10.2

Temple Grandin (b. 1947)



Temple Grandin is Professor of Animal Science at Colorado State University, but she is no ordinary professor. Diagnosed with

brain damage at the age of two, she is autistic, acts as a spokesperson for those with autism, and invented a 'hug box' or 'squeeze machine' to calm others. After a difficult and miserable time at school, she not only researched and wrote about autism but did a PhD on environmental enrichment for pigs. She believes that the autistic person's sense of being feared, dismissed, and threatened by everything gives her special insight into animals' experiences, noting that cattle are often disturbed by things most people don't even notice. Although redesigning slaughterhouses may not sound like a compassionate job, this is one way she has used her understanding of animal minds to improve their welfare. Her life was the subject of a 2010 biographical film, and a documentary about her was entitled 'The woman who thinks like a cow'.

mirrors (Bahrick, Moss, and Fadil, 1996). Children start referring to 'me' and then 'you' between eighteen months and two years of age, and around the same age start to manifest 'secondary emotions' like embarrassment and pride, suggesting that at this point they are starting to evaluate themselves in relation to the social world (Rochat, 2003).

A series of experiments by biologist Daniel Povinelli (2001) tested how children's self-recognition varies depending on whether they are looking at themselves in a mirror, a photograph, a video recording, or a live video. The child plays an unfamiliar game during which the experimenter praises the child by patting him or her on the head, and puts a big brightly coloured sticker there. Two- to three-year-olds had no difficulty recognising themselves (as 'me' or by using their name) in a video recorded three minutes earlier, but only 37% reached up to touch their head and find the sticker. When watching live video feedback, though, 62% reached up, and with a mirror 85% did, while with a photograph only 13% did. So, maybe recognising a 'present self' is easier than recognising a 'temporally extended self'. And putting it all together can be tricky: one 3-year-old said in response to questions, 'it's Jennifer' and 'it's a sticker', but then added, 'but why is she wearing my shirt?' (Povinelli, 2001, p. 81). Upbringing seems to make a difference too: 15- to 18-month-old infants from Scotland, Zambia, and Turkey, who interact with their mothers with varying amounts of physical or verbal contact, perform differently on tasks involving more or less autonomy: either recognising themselves in a mirror or recognising that their body is an obstacle to success in a task (Ross et al., 2017). And adults remain susceptible to the rubber-hand (Chapter 4) and body-swap (Chapter 17) illusions, and to alterations of self-recognition in many 'altered states of consciousness' (Chapters 13 and 15). So self-awareness is not all-or-nothing, even in humans.

But what about other animals? Are cats, dogs, or dolphins aware of themselves? Do they have a sense of 'I' as a conscious being observing the world? Would they be able to recognise themselves in a mirror?

Dogs and cats obviously cannot. Kittens will rush up to a mirror, look for the other kitten inside or run round the back to find it, and then quickly get bored. Many birds continue to treat their own image as a rival indefinitely, as do some fish. They clearly show no 'mirror self-recognition' (MSR). But what about our nearest relatives, the great apes?

Charles Darwin (1872) was the first to report the experiment. He put a mirror in front of two young orangutans at the zoo who, as far as he knew, had never seen a mirror before. He reported that they gazed at their own images in surprise, frequently moving and changing their point of view. They then approached close



FIGURE 10.9 • When we humans look in a mirror we recognise ourselves in the reflection, but which other animals can do this? Cats, dogs, and many other species treat their reflections as though they are another animal. Does mirror self-recognition imply self-consciousness?

and protruded their lips towards the image, as if to kiss it. Then they made all sorts of faces, pressed and rubbed the mirror, looked behind it, and finally became cross and refused to look any longer.

Sadly, we cannot tell whether these orangutans recognised themselves or not – whether they were looking at their own lips or trying to kiss another orangutan, for example. An attempt to find out more was not made until a hundred years later, when the comparative psychologist Gordon Gallup (1970) gave a mirror to a group of preadolescent chimpanzees. Initially they reacted as though they were seeing other chimpanzees, but after a few days they were using it to look inside their mouths or inspect other normally invisible parts of their bodies. Watching chimpanzees do this is certainly impressive. It seems obvious from the way they pick their teeth and make funny faces that they recognise themselves, but can we be sure?

To find out, Gallup anaesthetised these same animals and placed two red marks, one on an eyebrow ridge and one above the opposite ear. When they came round from the anaesthetic and looked in the mirror they saw the marks and tried to touch them, or rub them off, just as we would probably do. By counting the number of times the chimpanzees touched the marks compared with how many times they touched the same place on the unmarked side, Gallup could be sure that they did indeed see the reflection in the mirror as that of their own body.

Subsequently, many other species have been tested. Human children fail the test until they are somewhere between eighteen months and two years old. Chimpanzees vary a great deal, but generally do touch the spots. Of the three other species of great ape, orangutans and bonobos have been found to behave like the chimpanzees, but gorillas do not. Trying to give gorillas the benefit of the doubt, Gallup (1998) put marks on their wrist. They did indeed try to remove these marks but not the marks seen only in the mirror. The only gorilla to succeed has been Koko, a highly enculturated gorilla born in 1971 who has learned to communicate with humans using a modified version of American Sign Language. When asked what she saw in the mirror, she signed 'Me, Koko'. That Koko behaved so differently from other gorillas may seem surprising, but in fact it is well known that enculturated apes acquire many skills that their wild or captive conspecifics do not. Just what the relevant skills are in this case, though, we simply do not know.

In many similar tests, monkeys have shown no self-recognition, even though they use mirrors in other ways. For example, they can learn to reach things seen only in reflection, and will turn round towards someone they have seen in a mirror. Yet they do not pass the spot test. A possible reason is that while apes sometimes interpret eye contact as friendly, as humans do, most monkeys find it threatening and may not like looking in a mirror. Even so, placing mirrors obliquely to prevent eye contact does not seem to help.

MSR has sometimes been hailed as proof of a great divide in consciousness between us and the great apes versus all other animals, but this conclusion has been decisively overthrown. Dolphins and whales are extremely intelligent and communicative creatures, and some of them enjoy playing with mirrors. They have no hands to touch a spot, but there are other ways of measuring MSR. Working with two captive bottlenose dolphins who were used to mirrors, Diana Reiss

and Lori Marino (2001) marked them with either temporary black ink or just water on parts of their body they could not see. Both spent much more time twisting and turning in front of the mirror when the ink was used, in ways that would help them see the otherwise invisible marks.

Elephants are also highly intelligent, social animals with large brains, although radically different from apes and humans in their lifestyle and behaviour. Three Asian elephants were given large mirrors and not only passed the mark test but were found to go through the familiar stages of mirror use, progressing from social responses through physical inspections to testing the mirror with their own behaviour and finally apparently recognising themselves (Plotnik, de Waal, and Reiss, 2006).

Remarkably, some corvids do the same (Prior et al., 2008). Famously intelligent New Caledonian crows and jungle crows seem only to be able to use mirrors to explore the environment (Medina, Taylor, Hunt, and Gray, 2011). But when five European magpies were tested, they began by behaving as though another magpie were behind the mirror. Some were quite aggressive but then progressed to using the mirror in other ways, and three removed marks placed on their throats by looking in the mirror. What makes this so remarkable is that corvids' brains are quite different from those of great apes or elephants. The last common ancestor of mammals and birds was nearly 300 million years ago, and since then mammals have developed their layered cortex while birds developed a cluster of forebrain components. Bird brains are also tiny compared with ours, but their neurons are about twice as densely packed as those of mammals, and are especially dense in the forebrain (Olkowicz et al., 2016). Absolute brain size may also not be as important as size relative to body weight. In all animals that have passed the MSR test, the brain to body weight ratio is very high.

We still do not know for sure which species can and cannot recognise themselves in a mirror, but the test does seem to reveal an evolutionary convergence of abilities between radically different kinds of animals that are all sociable, intelligent, and capable of insight and imitation.

So what does MSR tell us about consciousness? It does not necessarily follow that because an animal can recognise its own body in a mirror, it has either self-awareness or a concept of self. For example, an ape might work out the contingencies between making movements and seeing effects in the mirror without concluding that the arm in the mirror is its own. Or a magpie might conclude that the mirror shows its own body without having any concept of itself as seen by others, or self as an agent or experiencer.

There is lively debate over this issue. Gallup (1998) is convinced that chimpanzees have not only MSR but also a concept of self and self-awareness. He even suggests that with this self-concept comes the beginnings of autobiographical memory and awareness of a personal past and future. In Damasio's (1999) terms (Chapter 16), this would imply extended consciousness and an autobiographical self as well as core consciousness.

Species	Brain Weight (g)	Body Weight (kg)	% Brain Weight × 1,000
European Magpie	5.8	0.19	31
African Grey Parrot	9.18	0.405	22.6
Pigeon	2.4	0.5	5
Human	1,350	65	21
Chimpanzee	440	52	8
Gorilla	406	207	2
Rhesus Monkey	68	6.6	1
Asian Elephant	7,500	4,700,000	1.6
Bottlenose Dolphin	1,600	170	9
Cat	25.6	3.3	8

FIGURE 10.10 • Encephalisation quotient (Cairó, 2011, p. 6). Can we tell an animal's intelligence from its brain size? Since brain size generally increases with body size methods have been devised for comparing relative brain size across species. The EQ, developed for comparing mammals, is a measure of how far the brain weight of a given species differs from that expected for an animal of its size. On this measure, although we do not have the largest brains in absolute terms, humans score highest. But no measure is perfect for assessing intelligence: for example, a thin and a fat person will have significantly different EQs. Even if we could accurately compare intelligence across (and between) species, would this tell us anything about consciousness?

'Can animals empathize? Yes.'

(Gallup, 1998)

'Can animals empathize? Maybe not.'

(Povinelli, 1998)

'Children and chimps and crows and octopuses are ultimately so interesting not because they are mini-mes, but because they are aliens – not because they are smart like us, but because they are smart in ways we haven't even considered.'

(Gopnik, 2016)

Povinelli (1998) agrees that chimpanzees have a concept of self, but not that they are aware of their own psychological states (for the counter-argument, see Gallup, 1998). He suggests that 'self-recognition in chimpanzees and human toddlers is based on a recognition of the self's behaviour not the self's psychological states' (p. 74). Others have argued that MSR is a mere by-product of a more general ability to collate and compare multiple mental models of the same thing – a skill demonstrated in search tasks and pretence as well as the ability to build and update expectations about their own physical appearance (Suddendorf and Butler, 2013). Even more sceptical is British psychologist Cecilia Heyes (1998), who agrees that chimps are capable of 'mirror-guided body inspection' (p. 102), but argues that they have no self-concept and no understanding of mental states. For her, 'mirror self-recognition' is the wrong way to describe the test; all it indicates, she says, is one form of mirror-guided exploration. But others have suggested that she takes such a hardline stance because she assumes that self-recognition is an all-or-nothing capacity which depends on being able to form second-order representations of oneself, rather than it being possible to recognise oneself in a more naïve way (Brandl, 2016).

KNOWING THAT OTHER MINDS EXIST

We humans have beliefs, desires, fears, and intentions, and we attribute them to others. That is, we have a 'theory of mind', or can 'mindread' or 'mentalise'.

Early theories of social cognition proposed that on the basis of what other people say, how they look, and what they do, I can infer their (unobservable) mental states, and on that basis I work out how to interact appropriately with them. This model, in which I construct a theory about the causes of other people's behaviour, is now often referred to by the shorthand of 'theory theory', and has more recently been challenged by 'simulation theory'.

In simulation theory, I understand others by running a simulation of their actions, as if performing them myself. The simulations can be conceived of as either conscious (or 'person-level', or 'explicit') or unconscious (or 'sub-personal', or neural, or 'implicit'). Either way, a simulation generates pretend versions of the other person's states in me, which allows me to grasp their thoughts, beliefs, and desires.

More recently still, 'interaction theory' has proposed that there is no need to posit any indirect, mental route to calculating someone else's 'inner state', whether by inference or simulation. On the contrary, I understand other people as I do myself, as an embodied agent in constant interaction with others, within a constraining and affording environment. To see someone smile is not a piece of evidence to feed into an inferential calculus or a mental simulator, but a direct experience of their pleasure or happiness. 'Accessing your thoughts, beliefs and desires thus becomes, for Interaction Theory, less a matter of reading your mind than attending to the world we already share' (Chesters, 2014, p. 71).

Possibly all three theories can help us understand how social cognition operates in different circumstances and responds to different challenges. Interaction theory certainly offers more scope for non-human interactions to qualify as fully fledged social cognition. All three, however, are compatible with Dennett's (1987) notion

of how readily we adopt 'the intentional stance'. That is, we understand other people's behaviour by treating them *as if* they have hopes, fears, desires, and intentions – just as we do with ourselves. The intentional stance is a very powerful tool for understanding, controlling, and predicting the world around us. It makes deception possible, as well as empathy.

Human babies are not born with these abilities. Sometime during their second year, they begin to follow another person's gaze to see what they are looking at, and to look at what is pointed at, rather than to the pointing finger. By the age of 3, they can talk about their own and others' desires and preferences. But at this age they cannot understand that someone else may not be able to see what they can see, or may have a false belief. This is the age at which a child playing hide-and-seek may hide her head under a pillow and shout 'come and find me'. Numerous experiments have shown that between the ages of 3 and 5 the various aspects of having a theory of mind develop.

In 1978 two psychologists David Premack and Guy Woodruff asked, 'Does the chimpanzee have a theory of mind?' The relevance of this question to us here is that if other animals do not have a theory of mind, and cannot attribute mental states to others or to themselves, it seems impossible that they could be conscious in the human sense. Mirror self-recognition is one aspect of this. Other relevant skills include the ability to understand what others can see or know, to deceive others, to empathise with others, and to imitate them.

Some monkeys give alarm calls to warn others of approaching danger. Calling is risky, and so it would be safest to call only when it could be useful. Yet many monkeys apparently call regardless of whether others have already seen the threat, or even whether there are any others around. The primatologists Dorothy Cheney and Robert Seyfarth carried out an experiment with a Japanese macaque mother. When put on the opposite side of a barrier from her infant, the mother gave the same number of alarm calls to an approaching predator whether or not her infant could see it. From this and many other studies, Cheney and Seyfarth (1990) concluded that monkeys do not have a theory of mind.

What about chimpanzees? Chimps will follow another's gaze, as though trying to see what the other is looking at. But this need not imply that they have a concept



CONCEPT 10

DECEPTION

- To deceive someone means to manipulate what they believe. A butterfly with a brilliant eye pattern on its wing deceives predators, as does a camouflaged stick insect, or a plover that feigns a broken wing to distract a predator away from its nest. In these cases, the camouflage or behaviour is genetically encoded, but human deception is rather different. You might deliberately try to convince someone that you didn't steal their chocolates or lose their book, or that you really do love them. You can only do this if you know that someone else can have a false belief.

This kind of social intelligence was largely underestimated until the 1980s, when Nicholas Humphrey argued that the need for social intelligence drove the increase in brain size among primates. With its emphasis on manipulation, deceit, and cunning, this became known as the 'Machiavellian Hypothesis' after Niccolò Machiavelli, the devious political advisor of sixteenth-century Italian princes (Whiten and Byrne, 1997).

Clearly humans are adept at deceit, but what about other primates? Many researchers working in the wild have reported fascinating stories (Byrne and Whiten, 1988). Monkeys and baboons will distract the attention of others in order to snatch food, or watch until others are fighting to grab an opportunity to mate with a receptive female. Rhesus monkeys may withhold their normal food calls so as to eat without sharing what they find, especially if they are very hungry or have found highly prized food. Swiss ethologist Hans Kummer watched for some twenty minutes while a female *Hamadryas* baboon gradually moved herself about two metres, while still sitting, until she was behind a rock where she began grooming a young male, behaviour that would be severely punished if the leading male saw her. Had she worked out what another baboon could and could not see?




FIGURE 10.11 • Deception and theory of mind are closely linked. Only a creature capable of attributing mental states to others would hope to get away with illicit activity by hiding behind a rock.

of what another chimp can see. They might have an evolved tendency to look where someone else is looking. To find out, careful experiments are needed.

Chimpanzees beg for food from humans and from each other. In an ingenious series of experiments, Povinelli (1998) and his colleagues used this behaviour to find out whether chimpanzees know what someone else can see. First they tested the chimps to make sure that they begged for food from an experimenter out of their reach, and did not beg for inedible items. Then two experimenters offered them food; one had a blindfold over her mouth and the other had one over her eyes. The chimps came into the lab, paused, and then begged for the food, but they were just as likely to gesture to the person who could not see them as the one who could. This was even true when one experimenter had a bucket over her head. Sometimes, when their begging failed to elicit any food, they begged again, as though puzzled at getting no response.

They seemed to pass one test: when one person turned her back the chimpanzees were less likely to gesture to her. However, when both experimenters sat with their backs to the apes and one looked back over her shoulder, the chimpanzees gestured randomly to both. They seemed oblivious to the fact that there is no point begging to someone who cannot see you. This is dramatically different from the behaviour of human children, who can understand this before they are three years old.

More recently, researchers have devised tests for theory of mind that do not require cooperation with humans. One experiment investigated whether subordinate chimps understand what their dominant counterparts know about food they are competing for. The researchers found that the subordinates can distinguish between cases where the dominant has or has not seen the food hidden or moved: subordinates will go for food unseen by the dominant competitor, but stay away when they know it has been seen (Hare, Call, and Tomasello, 2001).

Experiments have also distinguished between knowing about others' knowledge and about their beliefs. Two chimps take turns choosing from a row of buckets, some of them containing food. In the first condition (the knowledge-ignorance test), one chimp sees its competitor observing one of two pieces of food being hidden and then choosing one of three buckets. Can the chimp use its knowledge of what its competitor knows to determine which bucket might still contain food? In the second condition, the chimp sees its competitor misled by an experimenter pretending to put food in one bucket. Can it predict the competitor's choice (by identifying their false belief about the location of the food)? Six-year-old children pass both tests, but the chimps fail the false-belief test (Kaminski, Call and Tomasello, 2008). Some have speculated that apes can represent relations between agents and information that is true from their perspective, but cannot

'Unless one needs to discuss behaviour, or to catch a Hollywood spy, submentalising may be the smart option'

(Heyes, 2017, p. 2)

represent relations between agents and untrue states of the world (Martin and Santos, 2016).

In more recent experiments tracking the eye movements of various kinds of ape watching movies featuring humans, the apes have passed the false-belief test, too (Krupenye et al., 2016). Heyes (2017) argues, though, that what they were doing was not mentalising but 'submentalising': predicting behaviour using 'low-level, domain-general psychological processes' that did not evolve for reading others' minds (p. 1). But she stresses that this is not to belittle the apes: we use similar mechanisms a lot of the time, and 'Unless one needs to discuss behaviour, or to catch a Hollywood spy, submentalising may be the smart option' (p. 2).

Ravens have also been shown to differentiate between competitors who know where their food is hidden and those who do not (Bugnyar and Heinrich, 2005). In experiments with ravens hiding food, birds protected their caches and pilfered from others differently depending not just on whether they had seen other birds around them when the food was hidden, but also on whether or not there had been obstacles obstructing other birds' view of the hiding place.

Despite many ingenious experiments, Premack and Woodruff's question has still not been satisfactorily answered, and we still do not know which species, if any, have 'anything remotely resembling a "theory of mind"' (Penn and Povinelli, 2007). We have, however, learned that there are profound differences between the mental abilities of even closely related species, reminding us to take great care over any assumptions about consciousness.

IMITATION

Humans are 'the consummate imitative generalist', says psychologist Andrew Meltzoff (1988, p. 59). We imitate each other spontaneously and easily, and even infants can imitate sounds, body postures, and actions towards objects performed by adults. By 14 months of age toddlers seem to know when they are being imitated by adults and take pleasure in it (Meltzoff, 1996). As adults, we imitate far more than we may realise. We copy the body language of people we like and mirror their facial expressions when engrossed in conversation. In this way, imitation underlies the capacity for empathy. Variations in levels of empathy (whether measured as a personality trait, or as a response to environmental cues, cultural differences, or similarity between the imitator and the imitated person) have also been shown to correlate with the amount of imitation, especially if the person being imitated is attractive (Müller et al., 2013). It is perhaps because imitation seems so easy that we tend to think of it as a trivial skill and assume that other animals can do it as easily as we can. They cannot.

Nineteenth-century scientists like George Romanes and Charles Darwin assumed that dogs and cats learned by imitation, and that apes could 'ape', and they




FIGURE 10.12 • Do chimpanzees have a theory of mind? Can they understand what another person can and cannot see? In Povinelli's experiments, chimpanzees were just as likely to beg for food from an experimenter who had a bucket over her head as from one who could see.

WHAT IS IT LIKE TO BE THAT ANIMAL?

• S E C T I O N F O U R : E V O L U T I O N

told many stories of actions that looked like imitation. In 1898 the psychologist Edward Lee Thorndike defined imitation as 'learning to do an act from seeing it done', which captures the notion that to imitate means to learn something new by copying someone else. Over a century later, it is clear that this is far from trivial. The observing animal must not only watch the model, but remember what it has seen, and then convert that into actions of its own – even though these actions may look totally different from its own perspective. Computationally, this is a very complex task.

It is now clear that, with the exception of some birds and cetaceans imitating songs, there are very few species that can imitate. Even some of the classic cases turn out to be explicable in other ways. For example, in the 1920s in England, two varieties of small bird, blue tits and coal tits, were found to be pecking the foil tops of milk bottles left on doorsteps. Ethologists studied the way the habit started in a few places and then spread contagiously across the country. But this turned out not to require true imitation at all. It seems more likely that once one bird discovered the trick by trial and error, the jagged pecked tops attracted the attention of more birds who then associated the bottle with cream (Sherry and Galef, 1984). This is a form of social learning but not true imitation.

Even the famous Japanese macaques who learned to wash sweet potatoes in the sea may not, in fact, have learned by imitation. Young macaques follow their mothers about, and it may be that once one female learned the new skill, others followed her into the water and then, by accident, dropped their sweet potatoes and learned the trick of getting clean and salty sweet potatoes for themselves. This would fit with the fact that the whole troop learned only very slowly (Hirata, Watanabe, and Kawai, 2001). Young human children, with their avid delight in imitation, would learn such a skill in a few minutes rather than years.

There is clear evidence of culture in chimpanzees, in that different groups of chimpanzees have different ways of processing food, fishing for termites with sticks, or using leaves to soak up water, but there is ongoing controversy over how much these cultural skills are learned by true imitation rather than by other kinds of social learning (Heyes and Galef, 1996; Tomasello, 1999; Zentall, 2006). This is relevant to the question of whether other animals have memes or not, and what is needed for cultural evolution to take off.

Links can be seen in other primates between emotional connection and shared physical action. For example, emotional proximity (amount of grooming between two individuals) correlates with how contagious baboons' yawning is, regardless of spatial proximity (Palagi et al., 2009). And capuchin monkeys behave more sociably towards humans who imitate them (Paukner et al., 2009). Although yawning is a reflexive and stereotypical chain reaction rather than an imitated action, and the monkeys were here responding to imitation but not performing it themselves, findings like these suggest that behavioural matching with imitative qualities has important links with social relationships, and so with the fabric of what may contribute to conscious experiences.

Beyond the primate world, some whales and dolphins have local dialects in their songs, or signatures by which they recognise other individuals, and they copy songs back after hearing them (Reiss, 1998). There is also evidence that captive dolphins can imitate the actions of their human keepers, which is particularly

'fundamentally, deep down, chimps just don't "get it"'

(Pinker, 1994, p. 340)

interesting since their bodies are so very different. If imitation implies the capacity for empathy, then it is perhaps to these cetaceans that we should look for clues. Although we do not yet know how widespread imitation is, we must conclude that it is much rarer than most people realise, and is probably confined to only a few big-brained and intelligent species.

This may be important for understanding human evolution, because memes are defined as ‘that which is imitated’. This means that only a species capable of copying another’s behaviour can have memes and sustain a culture based on memetic evolution. One theory is that imitation, not introspection, Machiavellian intelligence, or the capacity for symbolic thought, set humans on a different evolutionary path from other great apes; it was memetic evolution that gave us big brains and language (Blackmore, 1999).

Imitation may be relevant to consciousness for another reason. If the concept of self is a memplex (Chapter 11), then it is the ability to imitate that gives humans not only language but a sense of self and hence self-consciousness.

LANGUAGE

The greatest divide of all may be that we have language and other species do not. Using true language means putting arbitrary symbols together in an unlimited number of ways, using grammatical rules, to convey different meanings. Humans are the only species known to do this. For some, this does not matter: ‘higher animals are obviously conscious’, says Searle (1997, p. 5); ‘consciousness reaches down into the animal kingdom’, says Metzinger (2009, p. 19). To others, it makes all the difference: ‘Perhaps the kind of mind you get when you add language to it is so different from the kind of mind you can have without language that calling them both minds is a mistake’ (Dennett, 1996b, p. 17).

If language makes human consciousness the way it is, then the consciousness of other creatures must be quite different from ours. If human consciousness and the concept of self are illusions created by language, then other creatures might be free of those illusions. Alternatively, you might argue that language makes little difference – that the heart of consciousness is about having sensory awareness, thinking, feeling emotions, and suffering (Feinberg and Mallatt, 2016). In that case, the divide between us and other creatures would not be so wide.

Children everywhere pick up the language around them with extraordinary speed and agility, without being specifically taught, and without being corrected for their mistakes. They have what is sometimes called a ‘language instinct’ (Pinker, 1994). From birth, infants respond more to human speech than to other sounds, and as early as 1 month they seem to be able to distinguish between different speech sounds. By 6 months, babies start to produce speech-like sounds themselves, forming proper words by 12–18 months, and then gradually developing the ability to form sentences with grammatical structure. This basic sequence is much the same across cultures, but variations in linguistic structure and cultural environment affect the rate and manner of acquisition, as well as the relations that develop between literal and figurative language use and patterns of thought.

Other animals certainly have complex methods of communication. For example, bees can communicate detailed information about the direction and distance

• SECTION FOUR : EVOLUTION

of a food source by dancing. Peacocks communicate how strong and beautiful they are by flashing their enormous tails. Vervet monkeys make several different alarm calls for different kinds of predator. But in all these cases, the meaning of the signals is fixed and new meanings cannot be made by altering or combining old ones.

Many attempts have been made to teach human language to other animals, in particular the other great apes. Early attempts failed because other apes do not have the vocal apparatus needed to make the right sounds. Realising this, in the 1960s, Allen and Beatrix Gardner tried teaching American Sign Language (ASL) to a young chimpanzee, Washoe, who lived with them and was treated like a human child. Washoe certainly learned many signs, but critics argued that she did not understand what the signs meant, that the experimenters were erroneously interpreting natural chimpanzee gestures as signs, and that she was not really acquiring true language (Terrace, 1987; Pinker, 1994).

Subsequently other chimpanzees also learned ASL, as did some gorillas, including Koko and her companion Michael, and an orangutan, Chantek. Koko and Michael learned to sign phrases more than eight signs long, with consistent grammatical structures. In an impressive display of the cognitive capacity known as 'conceptual blending', which underlies figurative uses of human language, they have also created new signs out of compounds of known ones: *stuck-metal* to mean *magnet*, for instance, or *insult-smell* for *garlic*. Like Washoe, Chantek was fostered by humans from a young age and learned hundreds of signs, but he did not learn

them as a child would, just by watching. His hands had to be moulded into the right shapes. Now over 40 years old, he understands much spoken English, and when Sue spent a day with him she concluded that he understood the crucial difference between such commands as 'put the stick on the blanket' or 'put the blanket on the stick', suggesting some understanding of grammar. Even so, his own sentences tend to be short and repetitive, and are mostly demands for food. Other apes have learned to communicate using magnetised plastic chips on a board, or modified computer keyboards.

Despite the real achievements of these apes, there remain glaring differences between their use of language and that of human children. While children show a great delight in naming things and telling other people about them, the apes seem mostly to use signs as a way of getting what they want (Terrace, 1987). As Pinker puts it, 'fundamentally, deep down, chimps just don't "get it"' (Pinker, 1994, p. 340).

Apes, it turns out, may not be the best choice of animals to teach human language to. Alex, an African grey parrot, learned to answer complex questions about the shape, colour, number, and material of objects shown to him. And unlike the apes, he could pronounce English words easily. Bottlenose dolphins have been given interactive underwater keyboards with which they can ask for playthings and answer questions (Reiss, 1998). They can also imitate artificial




FIGURE 10.13 • Chantek was brought up like a human child and taught American Sign Language from an early age. He was also trained to play 'Simon says', but although he can laboriously imitate some human actions he does not seem to take delight in imitation as human children do (Photo: Stuart Conway/Camera Press).

sounds made by the keyboard and then use the sounds spontaneously. It seems possible that dolphins will prove better language learners than many apes have been, and even that they have their own underwater language, representing the shapes of objects using the complex clicks and whistles by which they echolocate (Kassewitz et al., 2016). These speculations aside, it seems that we humans are alone in our spontaneous use of true language.

THE OCTOPUS

So what is it like to be an octopus? Octopuses are invertebrates; they are classified specifically as molluscs, along with animals like clams that do not even have brains. But octopuses can discriminate between objects based on size, shape, and brightness; they can learn the right path to a reward and how to retrieve a crab from a clear bottle sealed with a plug. They have a sleep–wake cycle, and blow water at floating objects in play. They have complex sensory receptors and nervous systems, with as many neurons as some vertebrates – but with more neurons located in the tentacles than in the brain, each tentacle containing a complex semi-autonomous neural network.

It is hard to say just how intelligent an octopus may be, but is intelligence even relevant to consciousness? We will come back to this question in Chapter 12, where we explore how artificial intelligence and artificial consciousness intersect.

Meanwhile, is the octopus conscious? You might answer ‘Yes’: every creature lives in its own world of experiences, however simple or primitive its senses might be. You might even answer that a single octopus tentacle is conscious. David Edelman and colleagues declare that ‘it is not likely that the question, “what is it like to be an octopus tentacle?” will ever be posed by any rational philosopher’ (2005, p. 178), but is there any good reason not to ask?

On the other hand, you might say ‘No’: the octopus lacks some critical ability without which there is no consciousness, such as intelligence, a self-concept, theory of mind, memes, or language. If you wanted to be really sceptical, you might say that it is just as impossible to answer the question, ‘What is it like to be my partner?’, as it is to answer the question, ‘What is it like to be an octopus?’ None of us can ever know what it is like to be any other creature, nor be sure that there *is* anything it is like to be any other creature. The furthest step down this radical line is that human consciousness is a grand illusion and there is nothing it is like to be us. In that case there would be no sense in



ACTIVITY 10.1

Lab choice

In a ‘Balloon debate’, every participant has to convince the others that they should not be thrown out to save the sinking hot-air balloon. In the lab debate, there is an equally difficult choice to be made between species.

Imagine that just one animal is going to be released from being tested on in a pharmaceutical laboratory and returned to the wild. Which species should it be?

Choose several different species and someone to defend each one, or let students pick their own favoured species. Each person is given a set length of time (e.g. 2 or 5 minutes) to make their case. Afterwards, the audience votes on which animal is released. If the choice proves easy, vote on which should go second and third.

This debate can be held without prior planning. Alternatively, ask students to prepare their case in advance. They might bring photos, videos, or other kinds of evidence. They might learn about the social and communicative skills of their chosen species, or about its intelligence, capacity for insight, memory, sensory systems, or pain behaviour. The aim is to explore the nature of animal suffering.

• SECTION FOUR : EVOLUTION

asking any 'what is it like to be...?' questions, whether of an octopus, a friend, or myself.

In the afterglow of the Big Bang, humans spread in waves across the universe, sprawling and brawling and breeding and dying and evolving. There were wars, there was love, there was life and death. Minds flowed together in great rivers of consciousness, or shattered in sparkling droplets. There was immortality to be had, of a sort, a continuity of identity through replication and confluence across billions upon billions of years.

(Stephen Baxter, *Manifold: Time*, 1999/2015, p. 3)

READING

Bloom, P. (2004). *Descartes' baby: How child development explains what makes us human*. London, Heinemann (pp. 189–227).

On our intuitive dualism, and what children think about death, magic, the supernatural, and God.

Brandl, J.L. (2016). The puzzle of mirror self-recognition. *Phenomenology and the Cognitive Sciences*, online first.

Recognising oneself in a mirror is not an all-or-nothing phenomenon, and the process of learning to do so varies between species including humans.

Burghardt, G.M., and Belkoff, M. (2009). Animal consciousness. In T. Bayne, A. Cleeremans, and P. Wilken (Eds), *The Oxford companion to consciousness* (pp. 39–53). Oxford: Oxford University Press.

Includes subsections on dolphins, great apes, and ravens, plus 'animal metacognition and consciousness'.

Dawkins, R. (1986). Explaining the very improbable. In R. Dawkins, *The blind watchmaker* (pp. 1–18). London: Longman.


How natural selection allows us to do this.

Key, B. (2016). Why fish do not feel pain. *Animal Sentience: An Interdisciplinary Journal on Animal Feeling*, 1(3), 39.

Arguments based on neuroanatomical considerations and the principle that structure determines function. Recommended commentaries include Morsella and Reyes, Gagliano.

Rochat, P. (2003). Five levels of self-awareness as they unfold early in life. *Consciousness and Cognition*, 12, 717–731.

Evidence for five stages in humans: progressing from confusion to self-consciousness from third- and first-person perspectives.



The function of consciousness

CHAPTER

ELEVEN

'the story of the emergence of consciousness seems to remain in medieval darkness'

(Dehaene, 2014, p. 7)

'Either we throw away the idea that consciousness evolved by natural selection, or else we have to find a function for it'

(Humphrey, 1987, p. 378)

CONSCIOUSNESS IN EVOLUTION

Evolutionary theory is especially good at answering 'why' questions. Why are leaves flat and green? So they can photosynthesise efficiently. Why do cats have fur? To keep them warm. Why do birds have wings? So they can fly. Why are we conscious? So we can. . .

It is easy to think that since humans are conscious, consciousness itself must have a function and be adaptive. Nicholas Humphrey makes this sound obvious: 'either we throw away the idea that consciousness evolved by natural selection, or else we have to find a function for it' (Humphrey, 1987, p. 378). He says 'we can take it for granted that – like every other specialized feature of living organisms – it has evolved because it confers selective advantage' (2011, p. 14). If he is right, we have to discover what that selective advantage was.

But the connection between consciousness and evolution may not be so simple.

Evolved traits are not necessarily adaptive traits, and there are other options. The field of evolutionary psychology can help us think more clearly about how and why the human mind has evolved the way it has, but its history has been a controversial one.

The principles of evolution apply as much to human beings as to slugs and beet-root plants, yet resistance to this idea has always been strong. At the end of *The Origin of Species*, Darwin suggested that 'Much light will be thrown on the origin

of man and his history' (1859, p. 488) and that psychology would find a secure foundation in biology. But it was many years before he discussed how, in *The Descent of Man* (1871). In the 1960s, Williams pointed out how difficult it is for 'people to imagine that an individual's role in evolution is entirely contained in its contribution to vital statistics [. . .], that the blind play of the genes could produce man' (Williams, 1966, p. 4).

This opposition reached its height with the publication, in 1975, of *Sociobiology: The New Synthesis*, in which biologist Edward O. Wilson explored the evolution of social behaviour, including that of human beings. For this he was abused and heckled, and even had water thrown over him when he lectured on the subject. Perhaps it is for such emotional reasons that the term 'sociobiology' is rarely used today, but many of its principles survive in the newer field of evolutionary psychology.

These two fields have much in common. For example, both have explored how human sexual behaviour and sexual preferences have evolved, whether there are sex differences in ability and aptitudes or just in socially created gender roles, and what the evolutionary roots are of aggression and altruism. Both try to account for, and assume there is such a thing as, human nature. Among the founders of evolutionary psychology, John Tooby and Leda Cosmides (2005) describe its goal as 'mapping universal human nature'. We are not a 'blank slate,' says psychologist Steven Pinker (2002). Nor are we noble savages corrupted by society, or beings imbued with a soul. Contrary to the currently dominant view in many intellectual circles, we are not capable of learning absolutely anything or escaping our evolved abilities and tendencies. We must learn to understand and accept human nature. Evolutionary psychology has been attacked for being reductionist, determinist, and adaptationist (e.g. Rose and Rose, 2000), but arguably these criticisms target a distorted idea of what the field actually is.

Unlike sociobiology, evolutionary psychology treats the human mind as a collection of specialised modules, or information processing machines, that evolved to solve particular problems – a view often caricatured as the 'Swiss army knife' view of the mind. Although we all share the same collection of evolved modules, each of us behaves in our own unique ways, depending on the genes we were born with and the environment in which we find ourselves. Sadly, few evolutionary psychologists have concerned themselves with consciousness, or asked whether there is a consciousness module, but Pinker lists some of the most troubling questions about consciousness, like whether your experience of red might be the same as mine of green, whether your visual system could be kept alive in a dish and have visual experiences, and whether beetles enjoy sex. He comes to the pithy conclusion, 'Beats the heck out of me!' (Pinker, 1997, p. 146).

Another difference is that whereas sociobiologists tended to treat most human traits as adaptations, evolutionary psychologists emphasise two reasons why they may not be. First, most of human evolution took place when our ancestors lived on the African savannah as hunters and gatherers. So we need to understand which traits would have been adaptive then, not which might be adaptive now (Barkow, Cosmides, and Tooby, 1992; Buss, 1999). So, for example, a taste

'An important feature of consciousness is that it seems to break the modularity of mind'

(Andrade, 2012, p. 596)

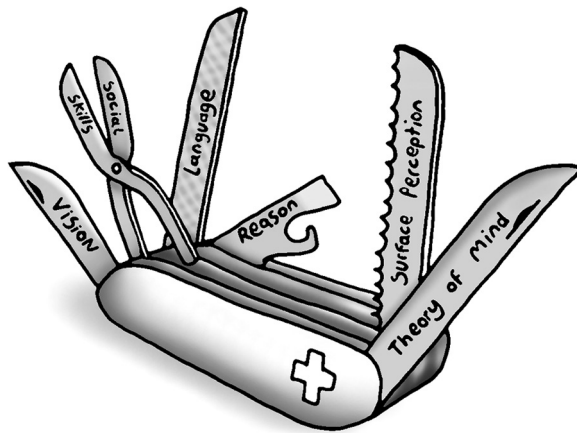


FIGURE 11.1 • In the Swiss Army knife caricature of the mind there is a special tool for every essential task. But how many modules are there? How much do they interact, and how specialised do they have to be? Is consciousness just one more blade on the knife?

for sugar was adaptive for a hunter-gatherer even though it leads to obesity and heart disease today; sickness and food cravings in pregnancy may have protected a foetus from poisons then, although well-fed women need different kinds of protection now; and superior spatial ability in males may have been adaptive when males were predominantly hunters and females were gatherers, even though we all have to get around vast buildings and cities today.

Second, evolutionary psychology emphasises the difference between the replication strategies of genes, and human strategies for gaining pleasure or success. As Pinker puts it,

almost everyone misunderstands the theory. Contrary to popular belief, the gene-centered theory of evolution does *not* imply that the point of all human striving is to spread our genes. [. . .] People don't selfishly spread their genes; genes selfishly spread themselves. They do it by the way they build our brains. By making us enjoy life, health, sex, friends, and children, the genes buy a lottery ticket for representation in the next generation, with odds that were favorable in the environment in which we evolved.

(1997, pp. 43–44)

In other words, we like good food and crave sex because people with those desires would, in the past, have been more successful at passing on their genes.

One of the results of our evolved nature is morality. Although genes are selfish, we are not – at least not always. Morality and consciousness have long been interlinked (Frith and Metzinger, 2016). Indeed, the words 'consciousness' and 'conscience' both come from the same Latin root *conscire* (*con-*, with, and *scire*, to know), via *conscius* (meaning having common knowledge with another) and *conscientia* (meaning moral conscience). This makes sense in that shared understanding underlies our ability both to understand ourselves and to empathise with others and so to care about them (Chapter 10). But does consciousness, in

the modern sense of subjectivity, play a role in moral decision-making? Is the capacity for moral action one of the adaptive functions of consciousness?

Here is a simple moral dilemma – and a true story. One night Sue’s son rang and told her that a publisher had offered him \$500 to use a photograph of his on a textbook cover. He was delighted but asked them to send him the fee in pounds instead of dollars. To his surprise, he received two cheques – one in pounds and one in dollars. He was thinking of tearing one up; should he?

What do you think he should do? And what advice do you think Sue, as his mother, should give him? Why?

This scenario involves a moral choice only because we have concepts of right and wrong, of fairness, of stealing, and of justice. We are not entirely selfish creatures and we do care about others and about behaving well ourselves. Where do these feelings come from? Some people believe that a God-given soul or spirit is the source of morality. Others think that moral decisions require consciousness, or even that one of the functions of consciousness is to guide morality. Others think that different societies impose different ethical concepts on their members and that morality is relative. Yet, the more we learn about the evolution of morality, the more clearly we can see its origins in our ancestors.

One factor is kin selection. All animals that care for their young, including humans, also care for other close relatives. This is because those relatives share some of their genes, and so aid to relatives is also aid to some of one’s own genes – depending on the closeness of the relationship. Another factor is reciprocal altruism, or doing good so that good will be done to you, which can be observed in many species, including vampire bats who share meals of blood and tiny cleaner fish that clean larger fish without being eaten by them, as well as chimpanzees and wolves. Often the favour has to be paid back at a later time, and this means that individual animals must be able to recognise each other and keep track of who has and has not reciprocated. They must then keep cooperating with the good sharers and punish the freeloaders, otherwise cheats would make successful sharing impossible. ‘Human beings [. . .] are uniquely good at reciprocal altruism’ (Matt Ridley, 1996, p. 84).

The evolution of reciprocal altruism is thought to have generated gratitude, sympathy, guilt, friendship, and trust, as well as moralistic aggression (the punishment of offenders) and the giving of gifts. Models derived from the mathematics of game theory have shown that certain types of behaviour, and certain mixtures of cheating and altruism, are more stable than others. In these and many other ways, we can understand how and why we humans have evolved the capacity for morality and our notions of fairness, trustworthiness, and virtue (Matt Ridley, 1996; Bloom, 2004; Hauser, 2006; Joyce, 2007).

Do we need consciousness on top of all these evolved tendencies in order to make truly moral decisions, or to think about such issues as aid to Africa, taxation and healthcare policy, abortion or assisted suicide? To stick with the simple dilemma above, if Sue’s son decides to tear up the cheque, why will he do so? Among possible answers are that his separate conscious mind intervened in the brain’s activity to make the moral decision, that his consciousness is an emergent property which can influence the decision through downward causation, that consciousness is an epiphenomenon and played no role at all in this or any other

• SECTION FOUR : EVOLUTION

moral decision, and/or that any power of consciousness is illusory and he only felt as though he consciously decided because of the way his brain functions.

These answers to the morality question all have very different implications for understanding the evolution of consciousness and why we are conscious at all.

Do you want to know what consciousness is for? Do you want to know the only real purpose it serves? Training wheels. You can't see both aspects of the Necker Cube at once, so it lets you focus on one and dismiss the other. That's a pretty half-assed way to parse reality. You're always better off looking at more than one side of anything. Go on, try. Defocus. It's the next logical step.

(Peter Watts, *Blindsight*, 2006, p. 302)

There are many people in the world who deny that consciousness did evolve. For example, some religious believers reject the very idea that humans evolved at all, despite the overwhelming evidence. Some Christians and Muslims believe that even if our bodies evolved like those of other animals, we alone have God-given souls and God alone can give us consciousness. Such souls are not dependent on the evolved body, so these ideas are thoroughly dualist.

A non-dualist can also deny that consciousness evolved by proposing that consciousness is fundamental to the universe and always existed, or that it is the power which drives evolution along, rather than being an evolved product itself. Yet few scientists, even those who believe in God, would want to 'throw away the idea that consciousness evolved by natural selection'. So let us accept the idea that consciousness evolved. Does this mean it must have a function?

At first sight, Humphrey's statements seem unexceptionable, and even look like a useful prescription for finding out why we are conscious. First, we find out what consciousness does, then we find out how that would have been useful for our ancestors' survival and reproduction. Then, hey presto, we have found out why consciousness evolved. But things are not so simple. Lurking within this apparently obvious statement are two, closely related, problems.

The first is this. When we asked 'What does consciousness do?' (Chapter 8), we found no easy answer. Indeed, a good case can be made for the idea that consciousness does nothing, or at least that it does nothing in its own right, separate from all the underlying processes that determine our behaviour. If consciousness does nothing, how can it have a function?

The second problem is related to the first. When we think about the evolution of consciousness, it seems easy to imagine that if things had turned out differently we might *not* have been conscious. The logic goes something like this.

I can see why intelligence has evolved because it is obviously useful.
I can see why memory, imagination, problem-solving, and thinking have evolved because they are all useful. So why didn't we evolve all these abilities 'in the dark'? There must have been some extra reason why we got consciousness *as well*.



PRACTICE 11.1

AM I CONSCIOUS NOW? DOES THIS AWARENESS HAVE A FUNCTION?

As many times as you can every day, ask yourself '**Am I conscious now?**' If you have been practising, you will know that asking this question seems to make you feel *more* conscious for a little while. Take this time to watch and wonder. Ask yourself, '**Does my awareness have any function of its own?**' Would my behaviour be any different without consciousness? If so, is this the kind of difference that natural selection could work on?

American philosopher Owen Flanagan uses this argument and takes it one step further to claim confidently that consciousness has no function. He says: 'Consciousness did not have to evolve. It is conceivable that evolutionary processes could have worked to build creatures as efficient and intelligent as we are, even more efficient and intelligent, without those creatures being subjects of experience' (Flanagan, 1992, p. 129). He calls this version of epiphenomenalism 'conscious inessentialism' and claims that its main thesis – that consciousness has no function – is both true and important.

Scottish psychologist Euan Macphail applies the same thinking to the *painfulness* of pain:

there does not in fact seem to be any need for the experience of either pleasure or pain. [...] What *additional* function does the pain serve that could not be served more simply by a direct link between signals from the classificatory system and the action systems?

(Macphail, 1998, p. 14)

You will probably have noticed something familiar about this argument. Yes, it is the zombie all over again (Chapter 2). If you believe in conscious inessentialism, then it follows that 'We might have been zombies. We are not. But it is notoriously difficult to explain why we are not' (Flanagan and Polger, 1995, p. 321). Or, to put it another way, 'it is hard to explain why evolution produced us instead of zombies' (Moody, 1995, p. 369). The idea that we could so easily have been zombies is so intuitively appealing that we must take it slowly and work out whether it really makes sense or not. In Chapter 2, we met some powerful reasons to reject the possibility of zombies, but for the sake of argument let us assume for now that zombies are (in principle, if not in fact) possible. This allows us to tell the imaginary tale of zombie evolution.

ZOMBIE EVOLUTION

As evolution proceeds, animals compete with each other to survive and reproduce, and traits like accurate perception, intelligence, and memory spread. One

'there does not in fact seem to be any need for the experience of either pleasure or pain'

(Macphail, 1998, p. 14)

'Consciousness did not have to evolve. [...] We might have been zombies.'

(Flanagan and Polger, 1995, p. 321)

creature becomes especially intelligent. There is, however, nothing it is like to be this creature, or any of the others. They are all zombies.

One day a strange mutation appears by chance in one of these creatures – the ‘consciousness mutation’. Instead of being a zombie, this creature is conscious. We can call it a ‘conscie’. Unlike all the other creatures, there *is* something it is like to be this first conscie. It suffers, it feels pain and joy, it experiences the qualia of colour and smell, sound and taste. The birth of the conscie is like Mary the colour scientist coming out of her black-and-white room for the first time.

Now what? Will this chance mutation prove adaptive, and the gene for consciousness spread rapidly through the population? Will the conscies outperform the zombies and wipe them out? Or will the two continue to co-exist in an evolutionarily stable mixture? Might planet earth even be like this today, with some of us being zombies and some of us being conscies? Indeed, might some famous philosophers be zombies while others are real-live-properly-conscious people (Lanier, 1995)?

These questions *seem* to make sense. But we must remember to stick to a clear definition of the zombie. The most common definition is that a zombie is a creature who is physically and behaviourally indistinguishable from a conscious human being. The *only* difference is that there is nothing it is like to be the zombie. So what happens?

Absolutely nothing happens. Natural selection cannot detect any difference between the zombies and the conscies. As Chalmers points out, ‘The process of natural selection cannot distinguish between me and my zombie twin’ (Chalmers, 1996, p. 120). They look the same and they act the same. They both do exactly the same thing in the same circumstances – *by definition* (if you argue that they don’t, you are cheating). If such a mutation were possible, then it would be entirely, and necessarily, neutral, and would make no difference at all to the way these creatures evolve.

This line of thought leads to an impasse. If we believe in the possibility of zombies, we find it natural to ask why evolution did not make us zombies. But then we find we cannot answer the question because (*on the definition of a zombie*) natural selection cannot distinguish between conscies and zombies.

This whole horrible problem is caused by the mis-imagination of zombies, says Dan Dennett. Zombies are preposterous, but by persistently underestimating their powers (making them unable to do things we think we need consciousness for), and hence breaking the rules of the definition, philosophers make them seem possible (Dennett, 1995c). If you imagine complex organisms evolving to avoid danger without experiencing pain, or intelligent self-monitoring zombies evolving without being conscious like us (i.e. zimboes, Chapter 2), you are like someone who is ignorant of chemistry saying he can imagine water that is not H₂O.

‘To see the fallacy’, says Dennett, ‘consider the parallel question about what the adaptive advantage of *health* is. Consider “health inessentialism”’ (1995c, p. 324). Suppose that swimming the English Channel or climbing Mount Everest could in principle be done by someone who wasn’t healthy at all. ‘So what is health *for*? Such a mystery!’ (p. 325). But this mystery only arises for someone who thinks that

**DOES BEING AWARE
NOW HAVE ANY
FUNCTION?**

you can remove health while leaving all the bodily powers and functions intact. In the case of health the fallacy is obvious, yet people keep making the mistake with consciousness. They imagine that it is possible to remove consciousness while leaving all the cognitive systems intact. 'Health isn't that sort of thing, and neither is consciousness' (p. 325).

Douglas Hofstadter (2007) imagines a very fancy car that might, or might not, come with a chrome ornament in the shape of a Flash Gordon rocketship. But consciousness is not an orderable 'extra feature' like this. You cannot order a car with a tiny engine and then say to the dealer, 'Also, please throw in *Racecar Power* for me' (or rather, you can order it, but it won't arrive). Nor does it make sense to order a car with a huge engine and then ask how much more you have to pay to get *Racecar Power* as well.

For Hofstadter, consciousness is like the power of a well-built car: it comes with good design. For Dennett, when you have given an evolutionary account of the talents of zimboes in monitoring their own (unconscious) informational states, you have done the job. There is *not in addition* something called consciousness that has effects *in its own right*. On this version of functionalism (Chapter 8) or on any version of illusionism (Chapter 3), any creatures that could carry out all the functions we do would necessarily be conscious like us.

'what is health for? Such a mystery!'

(Dennett, 1995c, p. 325)

We can now see that Humphrey was wrong. It is not true that 'Either we throw away the idea that consciousness evolved by natural selection, or else we have to find a function for it'. The alternative is to accept that consciousness is more like health or horsepower than an optional awareness module. If we do that, the mystery changes and so does the task of understanding the evolution of consciousness. The mystery becomes why consciousness seems to be a high-spec upgrade when it is not. The task is not only to explain how evolution produced humans with all their particular skills and abilities, but also why creatures with those skills and abilities are conscious, or are under the illusion that they are conscious, like us.

With this in mind, we can now see that there are four ways of approaching the evolution of consciousness (see Concept 11.1). If you believe in physically and behaviourally indistinguishable zombies, then it is forever a mystery why consciousness evolved, and you might as well give up. If you reject the possibility of zombies, you have three choices. Consciousness must be something separable from all the other skills and abilities we have evolved, in which case the task is to explain the function of consciousness, and how and why it evolved in its own right. Alternatively, consciousness necessarily comes about when those skills and abilities evolve, and the task is to explain why. Finally, maybe we are deluded about the nature of consciousness and are trying to explain the wrong thing entirely. Then, we have to ask why we evolved to be so easily deluded.

WHEN CONSCIOUSNESS EVOLVED

Asking why consciousness evolved also means asking when. It seems reasonable to suppose that a few billion years ago there was no consciousness on this planet and now there is, but how could consciousness (or awareness or subjectivity)



FIGURE 11.2 • What has left this rabbit? Is it a life force or *élan vital*? Now that we understand how life perpetuates itself this concept is not needed. Will the idea of consciousness go the same way?

evolve out of unconscious matter? William James, a pioneer of evolutionary psychology, explained the central problem.

The point which as evolutionists we are bound to hold fast to is that all the new forms of being that make their appearance are really nothing more than results of the redistribution of the original and

unchanging materials. The self-same atoms which, chaotically dispersed, made the nebula, now, jammed and temporarily caught in peculiar positions, form our brains; and the 'evolution' of the brains, if understood, would be simply the account of how the atoms came to be so caught and jammed. [...] But with the dawn of consciousness an entirely new nature seems to slip in.

(1890, i, p. 146)

FOUR WAYS OF THINKING ABOUT THE EVOLUTION OF CONSCIOUSNESS

1 Conscious inessentialism (epiphenomenalism)

Zombies are possible. In principle, there could be creatures that look and act exactly like us but are not conscious. Consciousness is separable from adaptive traits such as intelligence, language, memory, and problem-solving, but it makes no detectable difference (this is the definition of a zombie) and has no effects (this is epiphenomenalism). For this approach, the important (and mysterious) question is, 'Why did evolution produce conscious instead of zombies?'

2 Consciousness has an adaptive function

Zombies are not possible because having consciousness makes a difference. It is separable from evolved adaptive traits such as intelligence, language, memory, and problem-solving, and adds something



CONCEPT 11.1

James set himself the task of trying to understand how consciousness could 'slip in' without recourse to a mind-stuff, mind-dust, or soul. This is essentially the task we face today, but we should not confuse it with two other related questions. The first concerns when consciousness arises during human development. For example, is an unfertilised egg or a human foetus conscious? And if not, when does a baby or a child become conscious? The second (Chapter 10) concerns which creatures alive today are conscious. Answers to these may, or may not, help us with the question at issue here: when did consciousness first evolve?

Over this question there is strong disagreement. Some believe that its appearance was gradual, such as Susan Greenfield, who claims that 'consciousness is not all-or-none but comes in degrees', increasing like a dimmer switch with increasing brain size (Greenfield, 2000, p. 176). Others think quite the reverse. 'One thing of which we can be sure is that wherever and whenever in the animal kingdom consciousness has in fact emerged, it will not have been a gradual process' (Humphrey, 2002, p. 195).

Some place its arrival very early. For example, panpsychists believe that everything is conscious, although the consciousness of stones and streams is much simpler than that of slugs and sea lions. On this view, consciousness itself came long before biological evolution began, but the kind and complexity of consciousness might still have evolved. Some believe that life and consciousness are inseparable, so that as soon as living things appeared on earth, approximately four billion years ago, there

would have been consciousness. Some people equate consciousness with sensation, in which case it would have appeared with the first sense organs. The problems here concern defining sensation. For example, does the sunflower's ability to turn towards a source of light count as sensation and hence consciousness? Is the bacterium following a chemical gradient aware of the concentration it responds to?

Others think that consciousness requires a nervous system of some particular level of complexity or that it needs a brain, in which case consciousness would have appeared when these structures evolved. Defending 'the ancient origins of consciousness', Todd Feinberg and Jon Mallatt (2016) lay out what they call 'the defining features of consciousness' (p. 18). Being alive is not sufficient, nor is having a simple nervous system with reflexes. More complex neural hierarchies are needed that can create isomorphic representations – that is, representations which map features directly from the outside world onto the sensory system. This, they suggest, happened sometime between 560 and 520 million years ago, during the time of the Cambrian explosion, perhaps with a creature like *amphioxus*, a simple fish-like marine animal. 'The isomorphic visual images were processed by the expanding brain into mental images, which we propose marks the arrival of consciousness' (2016, p. 92). They can propose these 'defining features' as the criteria for consciousness (Chapter 10), but it is hard to see how their proposal could be tested, and others make equally specific and very different claims about the arrival of consciousness. Humphrey places it at 300 million years ago, and Bernard Baars (2012) ties it to the emergence of the mammalian brain around 200 million years ago.

Finally, there are those who believe that consciousness is a much more recent phenomenon, dating from the appearance of specialised social skills in our recent ancestors. Those skills include social perception, imitation, deception, theory of mind, and language.

The most recent origin for consciousness was suggested by American psychologist Julian Jaynes in his controversial book *The Origin of Consciousness in the Breakdown of the Bicameral Mind* (1976). Going back three thousand years to the earliest written records, he searched for clues to the presence or absence of a subjective conscious mind. The first text that allowed him accurate enough translation was the *Iliad*, an epic story of revenge, blood, and tears which describes events that probably occurred around 1230 BC and were written down around 900 or 850 BC. 'What is mind in the *Iliad*?' asks Jaynes. 'The answer is disturbingly interesting. There is in general no consciousness in the *Iliad*' (p. 69).

new. The important question is, 'What is the function of consciousness?', or 'What does consciousness do?'

3 Consciousness has no independent function

Zombies are not possible because any animal that could do everything we do would necessarily be conscious. Consciousness is not separable from evolved adaptive traits such as intelligence, language, memory, and problem-solving. The important question is, 'Why does consciousness necessarily come about in creatures that have evolved abilities like ours?'

(Note that functionalism falls into this category, but the term can seem confusing in this context. Functionalism claims that mental states are functional states, so explaining the functions performed also explains consciousness.)

4 Consciousness is illusory

Our ideas about consciousness are so confused that we fall for the zombic hunch, invent the hard problem, and worry about why consciousness evolved. The relevant question is, 'Why are creatures with abilities like ours so deluded about their own consciousness?'

'with the dawn of consciousness an entirely new nature seems to slip in'

(James, 1890, i, p. 146)

'consciousness is not all-or-none but comes in degrees'

(Greenfield, 2000, p. 176)

'it will not have been a gradual process'

(Humphrey, 2002, p. 195)



ACTIVITY 11.1

The sentience line

Is a stone conscious? Is a rose bush? Is a tadpole or a sheep? Is a baby? Are you? Where do you draw the line?

Gather together a collection of objects that you think span the range from definitely unconscious to definitely conscious. If you are doing this at home, you may have a pet to represent the animals, and house plants or a bunch of flowers for the plant kingdom. Indeed, you may be able to see enough examples just sitting in your own kitchen. Lay them out in front of you from the least to the most conscious and take a good look.

Doing this in class you may need to be more inventive, but having actual objects there forces people into making decisions and brings their arguments to life. You might ask people to bring in:

- 1 A stone or pebble
- 2 A weed from the garden, a houseplant, or a piece of fruit
- 3 A fly, spider, or woodlouse (put them back where you found them)
- 4 Tadpoles or pet fish
- 5 A thermometer
- 6 A phone
- 7 A human volunteer

Ask everyone to draw their own sentience line. Select the two people with the most extreme lines and ask them to defend their decisions against questions from the class. Does anyone move their line afterwards?




FIGURE 11.3 • The sentience line. Which of these do you think is conscious? Where do you draw the line?

*Deeply troubled he spoke to his own great-hearted spirit:
'Ah me! If I go now inside the wall and the gateway,
Poulydamas will be first to put a reproach upon me,
[...]*

*Now, since by my own recklessness I have ruined
my people,
I feel shame before the Trojans and the Trojan
women with trailing
robes, that someone who is less of a man than I will
say of me:*

*Hektor believed in his own strength and ruined
his people.' [...]*

(Homer, *The Iliad*, XXIII, ll. 13–15, 8th–7th century BC,
translated by R. Lattimore, 1951)

What he means is that there are no words for consciousness, nor for mental acts. Words that later come to mean 'mind' or 'soul' mean much more concrete things like blood or breath. And there is no word for will and no concept of free will. When the warriors act, they do so not from conscious reasons, motives, or plans but because the gods speak to them. In fact, the gods take the place of consciousness. This is why Jaynes describes these people's minds as 'bicameral' (meaning two-chambered). They were split. Actions were organised without consciousness, their motivations being heard as voices. We would now call these voices hallucinations, but they called them gods. So, 'Iliadic man did not have subjectivity as do we; he had no awareness of his awareness of the world, no internal mind-space to introspect upon' (p. 75). On Jaynes's view, our modern conception of consciousness as *subjectivity* describes something that is itself a recent invention. This view is similar to higher-order theories in defining consciousness in terms of 'awareness of awareness', and it also fits with an illusionist approach, since Jaynes thinks of consciousness as a 'learned cultural ability' (p. 380) and a 'metaphor-generated model of the world' (p. 66). When we locate consciousness inside our heads and even close our eyes to try to introspect on it better, we are subject to an illusion: 'In reality, consciousness has no location whatever except as we imagine it has' (p. 46). He offers one way of tracing the history of this illusion.

So, there is no consensus over when consciousness evolved, with theories placing its arrival anywhere between billions of years ago and only a few thousand years ago.

There is no consensus either over how or why consciousness evolved, but we are now ready to consider a selection from the many theories that try to answer these questions. In what follows we will be able to identify

which mystery they claim to be tackling and judge how well they succeed. We can bring some order to the chaos by thinking of the four options laid out in Concept 11.1. We have already explored the implausibility of zombie evolution and conscious inessentialism (consciousness exists but does not do anything). So we are left with three of our four options: possibly consciousness has a function in its own right; possibly, like health or horsepower, it just comes along with the whole; or possibly it is illusory (something exists, but not what we thought it was). There are theories of all three types.

'There is in general no consciousness in the Iliad'

(Jaynes, 1976, p. 69)

'Reflexes and simple motor programs; no need for consciousness there!'

(Feinberg and Mallatt, 2016, p. 62)

CONSCIOUSNESS HAS AN ADAPTIVE FUNCTION BIOLOGICAL FUNCTION

'Qualia are adaptive', claim Feinberg and Mallatt. 'Consciousness is a real, adaptive phenomenon that is of evolutionary survival value to the conscious organism' (2016, pp. 218, 217). They use the hard problem and the explanatory gap as markers for the arrival of sensory consciousness, assuming that both arose together and searching for 'the evolutionary origins of the gaps themselves' (p. 11). So, for them, the gap is a real phenomenon with ancient origins rather than being a recent problem invented by confused humans using language and philosophy. They claim that the hard problem can be solved with conventional biological principles and do this by assuming mental causation to be a real force that is visible to natural selection, although what this mental force does and how it works they do not explain.

'Consciousness is a supremely functional adaptation', proclaims Bernard Baars; and so he has to ask, 'how would you use consciousness, as such, to survive, eventually to pass on your genes?' (1997b, p. 157). His answer is that in our evolutionary past, consciousness would have saved us from danger – as in his example of escaping the full-sized angry bullock. But, as we saw in Chapter 8, he does not explain why it is 'consciousness, as such' – rather than having a global workspace architecture – which does the trick.

Max Velmans tries to answer by taking two perspectives. From a third-person perspective, he says, 'The same functions, operating to the same specification, could be performed by a nonconscious machine' (2000, p. 276), and so 'it is not obvious what the reproductive advantage of *experiencing* such information might be' (p. 277). His answer is that from a first-person perspective, life without consciousness would be like nothing, and 'there would be no *point* to survival' (p. 278). Yet he does not explain why life is not 'like nothing' or why there need be a *point* to our survival beyond the evolved instincts that keep us alive.

-
- Level 1: General biological features that apply to all living things
 - Life: embodiment and process*
 - System and self-organization*
 - Hierarchy, emergence, and constraint*
 - Teleonomy and adaptation*
 - Level 2: Reflexes that apply to animals with nervous systems
 - Rates and connectivity*
 - Level 3: Special neurobiological features that apply to animals with sensory consciousness.
 - Complex neural hierarchies; a brain*
 - Nested and non-nested hierarchical functions*
 - Neural hierarchies create isomorphic representations and mental images and/or affective states*
 - Neural hierarchies create unique-neural interactions*
 - Attention*
 - Sensory consciousness may be created by diverse neural architectures*
-

FIGURE 11.4 • The defining features of consciousness (from Feinberg and Mallatt, 2016, p. 18).

'Having established that consciousness is adaptive, we can now get down to "solving" the problem of subjectivity'

(Feinberg and Mallatt, 2016, p. 220)

'This consciousness that is myself of selves, that is everything, and yet nothing at all – what is it? And where did it come from? And why?'

(Jaynes, 1976, p. 1)

'Consciousness is a supremely functional adaptation'

(Baars, 1997b, p. 157)

Consciousness 'has a survival value in its own right', says Jeffrey Gray (2004). He rules out epiphenomenalism, arguing that language, science, and aesthetic appreciation would all be impossible without conscious experience, and that 'Whatever consciousness is, it is too important to be a mere accidental by-product of other biological forces' (p. 90). The good fit between our perceptions of the world and our actions in dealing with it cannot be an accident and must be under strong selection pressure. Of course, 'accidental by-product' and being 'under strong selection pressure' are not the only options, but by assuming they are, Gray can claim that this leaves us 'with the problem of identifying the causal effects of consciousness *in its own right*' (p. 90).

Gray rejects functionalism (the idea that mental states are functional states), claiming that synaesthesia (see Chapter 6) provides a counterexample. The colour qualia that synaesthetes experience, he claims, have no relationship to the word or number that triggers them and may even interfere with linguistic processing; and this, he says, is incompatible with functionalism. He argues that qualia are constructed by a chain of unconscious brain processes and are only correlated with or 'attached to' the functions that give rise to them; they are not the same as those functions.

Having rejected functionalism, and wanting to 'creep up' on the hard problem, Gray seeks 'the properties of qualia as such' (2004, p. 308). Conscious experience comes too late to affect rapid 'on-line' behaviour, he says, but slowly constructs the perceived world, smoothing out the moment-by-moment confusion to give a semi-permanent appearance. Since this seems to restrict any ongoing causal efficacy for consciousness, he concludes that it acts as a late error detector. An unconscious comparator predicts the next likely state of the world and compares this with the actual state. The results of the comparison then 'enter consciousness'. The hippocampal system is the neural substrate for the comparator system and 'so provides conscious experience with its evolutionary survival value' (p. 317). This is related to the earlier proposal of British ethologist John Crook that 'Consciousness of the world comprises a form of re-presentation of the current perceptual input on a mental screen so that a constant awareness of a monitoring process arises' (Crook, 1980, p. 35).

These theories give consciousness its own survival value and function, but do not explain why such monitoring or error detection requires or entails subjective experience or 'creates qualia' when other brain processes do not. Cartesian materialism is implicit in phrases like 'mental screen' and 'entering consciousness', and there remains a magic difference between brain processes with qualia 'attached' and those without. As Gray admits, he may have crept up on the hard problem, but he has neither explained it away nor solved it.



PROFILE 11.1

Nicholas Humphrey (b. 1943)



As a PhD student in Cambridge, Nicholas Humphrey discovered, almost by accident, that monkeys can still see after their visual cortex has been removed (the phenomenon later known

as blindsight). In 1971, during several months at Dian Fossey's gorilla research centre in Rwanda, he began to focus on the evolution of social intelligence, leading to the idea that human beings are 'natural psychologists' who use introspection to model the minds of others. He convinced Richard Dawkins that memes are living structures; made a 1980s TV series called *The Inner Eye*; and has long worked for the cause of nuclear disarmament. Returning to Cambridge in 1990 after three years with Dan Dennett at Tufts, he worked on radically new ideas about the nature of sensation and qualia, arguing that sensations are a form of 'bodily expression'. In *Soul Dust: The Magic of Consciousness*, he claims that phenomenal consciousness is a 'magical mystery show' designed by natural selection to have seemingly 'out-of-this-world' properties that make us feel special and transcendent.

Another set of theories also gives consciousness its own survival value, but for social rather than individual reasons.

SOCIAL FUNCTION

Once upon a time there were animals ancestral to man who were not conscious. That is not to say that these animals lacked brains. They were no doubt percipient, intelligent, complexly motivated creatures, whose internal control mechanisms were in many respects the equals of our own. But it is to say that they had no way of looking in upon the mechanism. They had clever brains, but blank minds. [They] . . . went about their lives, deeply ignorant of an inner explanation for their own behaviour.

(Humphrey, 1983, pp. 48–49)

So begins Humphrey's 'Just-So Story' of the evolution of consciousness: a story to explain how and why we humans became conscious.

Humphrey describes his own surprise and pleasure at coming across Wittgenstein and behaviourism, and discovering the 'naughty idea' that human consciousness might be useless. 'But it is a naughty idea which has, I think, had a good run, and now should be dismissed' (1987, p. 378). Consciousness must make a difference, he concludes, or else it would not have evolved.

Developing his theory in the 1980s, he treated consciousness as an 'emergent property'. Emergent properties are those, like wetness, hardness, or the weather, which are properties of a combination of things, but not of those things alone. So, for example, the wetness of water is not a property of either hydrogen or oxygen but emerges when the two form molecules of H₂O. Humphrey also describes consciousness as a 'surface feature', an emergent property that arises out of the combined action of the brain's parts, and on which natural selection can act. For example, the insulating properties of fur on an animal's body are a surface feature of hairy skin, and they are visible to natural selection, and have evolved for obvious reasons: the warmth is what matters to the animal's survival. So why did consciousness evolve?

Humphrey's answer is that the function of consciousness is social. Like our close relatives the chimpanzees, we live in highly complex social groups, and like them, our ancestors must have made friends and enemies, formed and broken alliances, judged who was trustworthy or not, and so needed the skills of understanding, predicting, and manipulating the behaviour of others in their group. In other words, they became 'natural psychologists'.

Rather than just watching others and noting the consequences, as a behaviourist might, imagine what would happen if one of these ancestral creatures could watch itself. Imagine that early hominid Suzy notices that ferocious Mick has a large piece of food and that her friend Sally is close by, obviously hoping to get some. Should Suzy join in and help Sally snatch it? Should she distract Mick by grooming him so that Sally can get it? If she does, will Sally share the food with her afterwards? By asking, 'What would I do in the circumstances?', Suzy the natural psychologist can make a better decision.

• SECTION FOUR : EVOLUTION

This is what we humans do, argues Humphrey, quoting from Descartes' contemporary Thomas Hobbes, who said:

Whosoever looketh into himself and considereth what he doth when he does think, opine, reason, hope, fear &c. and upon what grounds, he shall thereby read and know what are the thoughts and passions of all other men upon the like occasions.

(Hobbes, 1648/1946, in Humphrey, 1987, p. 381)

So Humphrey proposes that natural selection favoured self-reflexive insight. 'Now imagine that a new form of sense organ evolves, an "inner eye", whose field of view is not the outside world *but the brain itself*' (Humphrey, 1987, p. 379; 2002, pp. 74–75). In similar vein, Dawkins speculates that 'Perhaps consciousness arises when the brain's simulation of the world becomes so complete that it must include a model of itself' (Dawkins, 1976, p. 59).

The problem with both these ideas is that the picture we have of ourselves (if picture is even the right word) is not of glial cells, neurons, synapses, or brain activity, but of a person. Indeed, most people on the planet know nothing of what their brain would look like if they could see it. So perhaps the self-description needs to be understood in some other way, such as in Graziano's attention schema (Chapter 7) or Metzinger's self-model (Chapter 16), or in computational terms (Chapter 12), rather than as a model of the brain itself.

Humphrey describes the picture as a user-friendly description, designed to tell us as much as we need to know in a way we can understand. It allows us to see our own brain states *as* conscious states of mind. This, claims Humphrey, is what consciousness amounts to. It is a self-reflexive loop, rather like Hofstadter's strange 'I' loop (Chapter 16), whose function is to give human beings an effective tool for doing natural psychology.

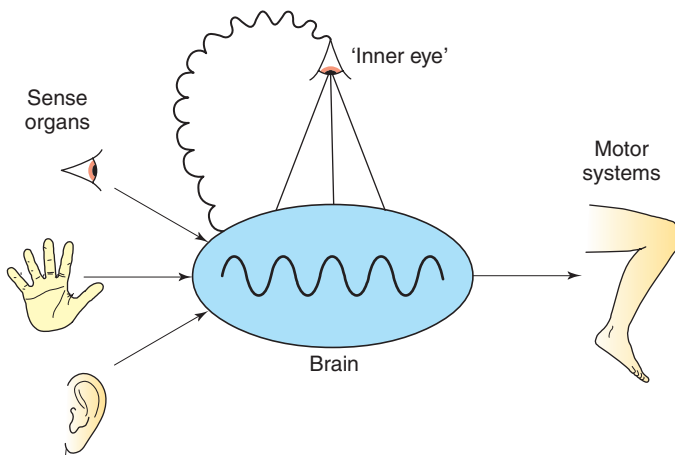


FIGURE 11.5 • According to Humphrey's earlier views, consciousness arose when a new form of sense organ evolved, an 'inner eye' whose field of view was not the outside world but the brain itself (Humphrey, 2002, p. 75; 1986, p. 70).

Other theories build directly on Humphrey's early ideas. For example, British archaeologist Steven Mithen (1996) agrees that consciousness has a social function and that chimpanzees probably have conscious awareness of their own minds. But he argues that if Humphrey is right, this awareness should extend only to thoughts about social interaction. Yet we humans seem to be conscious of all sorts of other things. It is this broadening of awareness that he sees as critical in the creation of the modern human mind.

Mithen likens our minds to a vast cathedral with many smaller chapels. In early hominid evolution, separate abilities evolved largely cut off from each other, like the modules in the Swiss army knife analogy. In *Homo habilis*, and even in the later Neanderthals, social intelligence was isolated from tool-making or interacting with the natural world: 'consciousness was firmly trapped within the thick and heavy chapel walls of social intelligence – it could not be "heard" in the rest of the cathedral except in a heavily muffled form' (Mithen, 1996, p. 147). These creatures, he supposes, had an ephemeral kind of consciousness with no

introspection about their tool-making or foraging, but with increasing cognitive fluidity the doors between the chapels opened, and the truly modern human mind evolved, coinciding with the cultural explosion of 60,000–30,000 years ago. By then our ancestors had already evolved big brains and language and were physically similar to us.

For Mithen, language evolved to substitute for grooming as the size of hominid social groups increased (Dunbar, 1996). On this theory, language was originally used only for talking about social matters, and even today the major topics of conversation between both men and women can be classified as ‘gossip’: that is, people talk about who said what to whom, who likes whom, and their own and others’ status and relationships (Dunbar, 1996). But once language had evolved, it could be used for other purposes, providing selection pressure to extend its use to talk about other matters such as hunting, foraging, and the physical world. This, argues Mithen, opened up the chapels of the mind. We have now lost our Swiss-army-knife minds and are conscious of much more than the social world that gave rise to awareness in the first place.

Other scientific theories also relate consciousness to our capacity for symbolic thought, such as Terrence Deacon’s (1997) theory of how the coevolution of the brain and language gave rise to the ‘symbolic species’, and Merlin Donald’s (2001) theory of the coevolution of human brains, culture, and cognition. This association with symbolic thought goes back at least to the ‘symbolic interactionism’ of the American philosopher and social psychologist George Herbert Mead. Mead argued that while other animals may be conscious, only humans have become self-conscious, and this self-consciousness is built up first from gestures and other nonsymbolic interactions and finally from the symbolic interactions made possible by language. For Mead, as for Russian psychologist Lev Vygotsky, consciousness came late in evolution and is fundamentally a social, not an individual, construction.

An interesting implication of these social theories is that only intelligent and highly social creatures can be conscious. These might include the other great apes, and possibly elephants, wolves, and dolphins, but most creatures throughout evolutionary history, and most alive today, would not be conscious at all.

One objection to Humphrey’s view is that introspection is unreliable. Even if we set aside the misleading metaphor of inner vision, the activity of introspection is responsible, for instance, for ‘convincing some people that their decisions have a kind of freedom that is incompatible with physical causation, or giving the impression

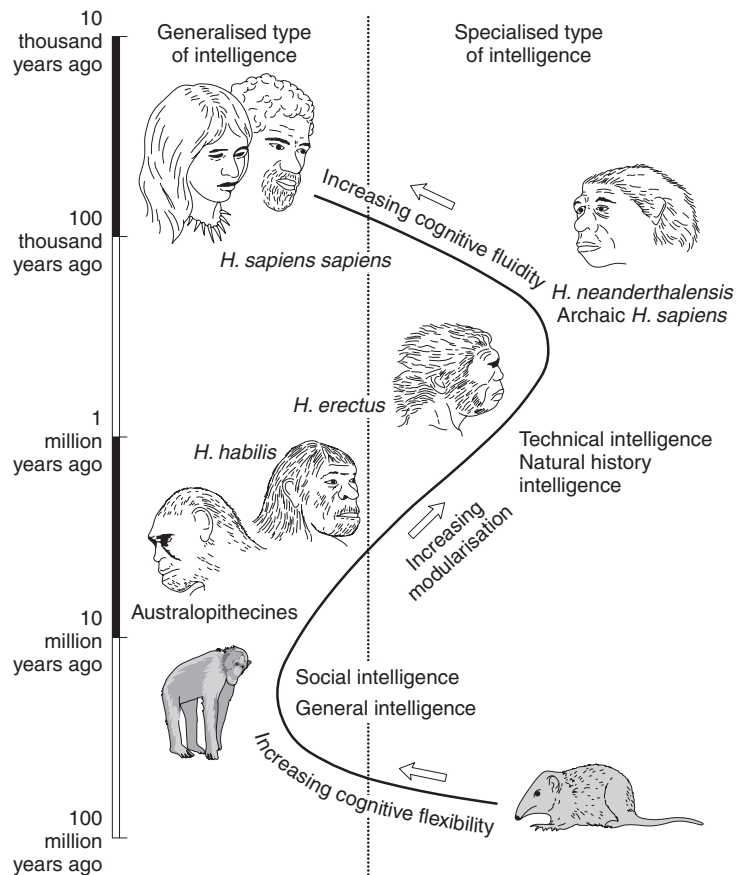


FIGURE 11.6 • Mithen suggests that during the evolution of the mind, selective advantages have oscillated between favouring specialised, hard-wired, or modularised intelligence and favouring general intelligence (Mithen, 1996, p. 211).

EXPERIENCE	INTROSPECTIVE MESSAGE	SURVIVAL VALUE
Pain	Unpleasant and to be avoided	Minimises injuries
Love	Desire for lifelong attachment, feelings of unbounded admiration, etc.	Propagation of the human species
Redness	Attribute of a physical object	Ability to communicate about this attribute

Introspection on our experiences does not directly tell us their survival value

FIGURE 11.7 • According to Barlow, introspection on our experiences does not accurately reflect their survival value (Barlow, 1987, p. 364).

that their visual field is filled with uniformly detailed information, or persuading them that they understand things they do not (Sloman and Chrisley, 2003, pp. 137–138). We might think we would be better off without pain, but in fact those few unfortunate individuals who cannot feel pain constantly damage themselves. Then there is that old chestnut, the redness of red – the ‘raw feel’, the quale. This we get completely wrong, says British physiologist Horace Barlow. When we say ‘This apple is red’, we may, from introspection, think that the raw sensation of red comes first, when in fact much computation is required, and the way we experience red depends on our whole history of seeing red objects and talking about them. Barlow argues that ‘the sensation of redness is merely preparing you to communicate the fact that something is red; this is another case where introspection is misleading, for redness is a carefully cooked product and is never as raw as it seems’ (1987, p. 372). This is reminiscent of James’s claim that ‘No one ever had a simple sensation by itself’ (1890, i, p. 224). Barlow concludes that consciousness is a social product derived from communication and cannot be explained by introspection.

Another objection is that Humphrey’s notion of consciousness is dualist, that the inner eye is a ghost in the machine or an audience of one in its Cartesian theatre, but he makes clear that this is not, in spite of the diagram, what he means. Rather, the inner eye is an aspect of the way the human brain functions. But how can a brain describe itself? Who is the observer inside the brain, and doesn’t this lead to an infinite regress of ever more observers? Humphrey says not. Consciousness, he says, is a feature not of the whole brain but only of the added self-reflexive loop whose output is part of its own input. No regress is implied, he claims. Yet he does admit there is a problem. ‘Why this particular arrangement should have what we might call the “transcendent”, “other-worldly” qualities of consciousness I do not know’ (2002, p. 75).

Do Humphrey and Mithen really see consciousness as *itself* having a function that is acted on by natural selection (Type 2 in Concept 11.1), or do they try to explain why any creature capable of introspection or self-reflective insight must inevitably be conscious (Type 3)?

The answer appears to be the former. Both Humphrey and Mithen describe consciousness as an emergent property with specific functions on which natural selection can act, such as ‘giving the subject a picture of his own brain activity’ (Humphrey, 2002, p. 76).

This may leave us with a fundamental doubt. Is consciousness really the kind of thing that can *be* a surface feature or an emergent property, like fur or wetness or intelligence? As ever, we must remember that consciousness means subjective experience, or ‘what it is like to be’. So the question for these theories is, does natural selection act on how it *feels* to introspect or on the behavioural consequences of introspection? If you decide the latter, then the subjective experience has no evolutionary function in its own right. Both its existence and the reason why it evolved remain unexplained.

Interestingly, Humphrey’s later work tries to avoid this problem and so belongs in the next section.

NO INDEPENDENT FUNCTION FOR CONSCIOUSNESS

An alternative approach is to deny consciousness a separate function of its own and ask instead how creatures built like us come to be conscious – much as we might ask how an animal comes to be alive or to be healthy, or how a car gets its horsepower.

Many kinds of theory fit this general approach. Perhaps the most extreme version is eliminative materialism. The Churchlands, for example, argue that once we understand the evolution of human behaviour, skills, and abilities, the whole idea of consciousness will just slip away, as did the idea of the ‘life force’ or ‘phlogiston’. In other words, there is no independent consciousness, and we need not ask how it evolved.

More common are those theories which deny consciousness a separate function without eliminating it. For example, psychologists Peter Halligan and David Oakley suggest that ‘It is our capacity to tell others of the contents of our consciousness that confers the evolutionary advantage – not the experience of consciousness itself’ (2015, p. 27). This, they say, is because communication about consciousness helps us predict the behaviour of others and respond to social influences. This is similar to Humphrey’s and Barlow’s social theories, except that it is communication, not subjectivity itself, that natural selection favours.

Many scientists working on human evolution avoid the tricky topic of consciousness, but implicitly adopt a functionalist position. Although some philosophers argue that functionalism cannot account for subjectivity at all (Chapter 8), others equate subjectivity with such functions as social interaction, language, or problem-solving. In this view, consciousness is not separate from those functions and so cannot have causal properties or functions in its own right (one of the reasons why the term ‘functionalism’ can be confusing). Explaining how the functional organisation of the brain and the rest of the body came about is all that is required.

Humphrey’s ideas moved on from his earlier work on ‘nature’s psychologists’ to provide a different take on how ‘the mind–brain identity equation for sensations *could* work’ (2006, p. 98). Having begun to sketch an action-based or enactive theory (Humphrey, 1992), he now tells a new just-so story (2006), starting with an amoeba-like creature that reacts to chemicals or vibrations around it by wriggling towards or away from them. At first, the responses are purely local, but soon they become linked into some kind of nervous system for more effective action. As sense organs evolve, the creatures react with more complex wriggles until the time comes when they need to make internal representations of their world. And here is Humphrey’s key point – they can do this by using their own reactions to the outside world, via the command signals they already issue. So this creature learns about the outside world and its own feelings ‘by the simple trick of *monitoring what it itself is doing about it*’ (2006, p. 87).

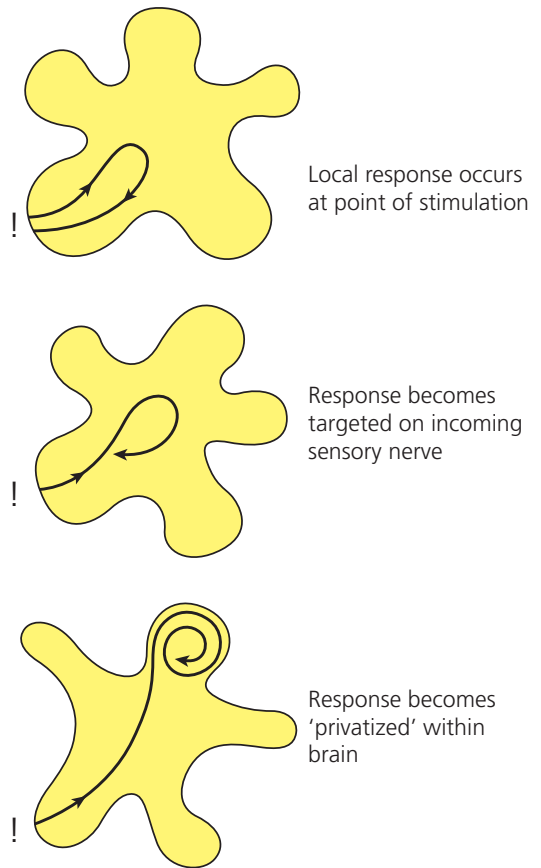


FIGURE 11.8 • The ‘privatisation’ of sensation (Humphrey, 2002, p. 112).

‘to sense the presence of red light, [the animal] monitors its signals for wriggling redly’

(Humphrey, 2006, p. 90)




FIGURE 11.9 • Humphrey (2011) likens consciousness to visual illusions such as these Penrose stairs.

'it is our capacity to tell others of the contents of our consciousness that confers the evolutionary advantage'

(Halligan and Oakley, 2015, p. 27)

'[Consciousness] is a magic show that you stage for yourself inside your own head'

(Humphrey, 2011, p. 199)

'illusionism [...] should be considered the front runner'

(Dennett, 2016, p. 65)

This, Humphrey argues, is how sensations came about. The response circuits were then privatised and internalised, and perceptions developed in a separate stream, but sensations still bear the hallmark of their origins in bodily expression. This changes the problem of consciousness by putting sensation on the side of agency rather than reception, and this is how Humphrey (2000) wants to solve the mind–body problem. Whether or not you agree that turning sensations into actions really does solve the mind–body problem, such enactive or sensorimotor theories do not need to ask how or why consciousness evolved in its own right because consciousness is not something separate from action.

This is also true of predictive-processing theories of consciousness (e.g. Seth et al., 2012), which treat consciousness as the result of top-down predictions generated throughout the course of our brain–body–world interactions. These probability-based predictions are shaped by prior knowledge accumulated over the lifespan. They take the form of interoceptive signals that, when they successfully match sensory inputs, yield conscious experience. In these respects, predictive processing can be seen as extending and adapting Humphrey's ideas by introducing the concept of Bayesian probability in a widely distributed system.

Humphrey goes on to explain why consciousness seems to matter to us so much: *'it is its function to matter'* and to seem mysterious and other-worldly (2006, p. 131). Ancestors who believed in a mysterious consciousness and an unworldly self would have taken themselves more seriously and placed more value on their own and others' lives. This is why belief in mind–body duality evolved.

So, on his 'reductionist theory of what phenomenal consciousness *is* – it is a magic show that you stage for yourself inside your own head' (2011, pp. 198–199). This sounds very much as though Humphrey is calling consciousness an illusion. He likens our 'magical mystery show' to such visual illusions as the impossible triangle or the Penrose stairs made famous by M.C. Escher's paintings of impossible stairways, calling it an 'ipsundrum': an 'illusion generating inner creation in response to sensory stimulation'.

Yet, in the end, Humphrey (2016) denies illusionism, and his denial keeps his theory firmly in this section. If 'sensations are representations of something *we do*' (2016, p. 117), internalising responses to incoming stimulation and making sense of them, then sensations have real effects in the world. Natural selection would operate on these effects and eventually lead to the kinds of minds we have today. Consciousness may be a magic show, but its effects are real enough.

THE EVOLUTION OF ILLUSION

Others disagree. Our last possibility is that phenomenal consciousness, as usually conceived, is illusory (Chapter 3). Illusions themselves may have effects, but experiences do not have phenomenal or 'what-it's-like-to-be' properties, and 'consciousness itself' does not exist. So natural selection has nothing to work on. Just as Frankish (2016b) has argued for replacing the hard problem with the illusion problem, so we should be replacing the question 'how did consciousness evolve?' with 'how did the *illusion* of consciousness evolve?'

We have already met one version of this in Dennett's zimbo (Chapter 2). The zimbo is a self-monitoring zombie, and because it can monitor its own internal states, it ends up speaking like we do about, for example, its thoughts, imaginings, and intentions. It believes it has phenomenal consciousness even if it doesn't. This is clearly an illusionist description and indeed Dennett goes so far as to call illusionism 'the obvious default theory of consciousness' (2016, p. 65).

Another version is Chris Frith's view that 'our brain creates the illusion that our own mental world is isolated and private' (2007, p. 17). On such theories, consciousness is not something extra with functions of its own. It is not a new emergent property on which natural selection can act. What has evolved is our capacity for thought and language, and our intuitions about ourselves, which may themselves be useful even though they also lead us astray. For example, being a dualist may have advantages even if dualism is false. So natural selection acts on the ability to think, talk, and monitor internal states, and the result is what we call a conscious creature. To the extent that such a creature believes something else about its own consciousness, it is suffering from an illusion.

Guy Claxton argues that consciousness started out as a rare phenomenon of super-alertness, a marvellous mechanism for spotting and responding to basic emergencies.

It did not emerge 'for a purpose'. It came along with the developing ability of the brain to create these transient states of 'super-activation' as a useless by-product, of no more functional interest than the colour of the liver, or the fact that the sea, under certain conditions, bunches up, rolls over and turns white.

(1994, p. 133)

'It came along [...] as a useless by-product'

(Claxton, 1994, p. 133)

The sad thing, says Claxton, is that while consciousness began as a rarity, we now live in an almost perpetual state of low-grade emergency, and it has become 'a mechanism for constructing dubious stories whose purpose is to defend a superfluous and inaccurate sense of self' (p. 150). This suggests the interesting conclusion that if we could learn to lead calmer and more mindful lives, our consciousness might not just change but even dissolve away (Chapter 18).

Our last theory, Graziano's attention schema theory, is hard to classify, because Graziano says that it has a lot in common with illusionism and belongs in the same category. But he, like Humphrey, balks at actually using the term 'illusionism' because it risks creating confusion and unwarranted backlash (2016, p. 112). On this theory, the brain doesn't just use attention, it constructs an internal model of it. This model, the attention schema, first evolved as a simple model of the organism's own state of attention. From there it evolved to modelling the attentional states of others and thereby predicting, understanding, and relating to them more effectively (see also Graziano and Kastner, 2011). This development was adaptive for three main reasons: integrating information, allowing increasingly efficient control of attention, and improving social skills. So the schema does exist and is firmly rooted in evolved brain mechanisms but is never literally accurate. Consciousness is what the internal model depicts, and what it depicts is a caricature, 'a cartoonish, somewhat inaccurate model of something real' (Graziano, 2013). Nonetheless, this is not

'To call consciousness an illusion risks confusion and unwarranted backlash'

(Graziano, 2016, p. 112)




FIGURE 11.10 • Tower of Generate-and-Test.

a brain error, but a useful and efficient adaptation: consciousness has adaptive functions (both biological and social), even if we could also describe it as, in one sense, illusory.

In fully illusionist theories, there is no need to explain how ‘consciousness itself’ or phenomenal experience evolved, because they did not. What evolved was our propensity to mischaracterise our own minds, creating the illusion of duality and inventing the hard problem. The nature, function, and origin of these illusions differs between the theories, but they all assume that evolution means biological evolution based on genes. There is an alternative and broader view of evolution based on what Dawkins (1976) calls ‘Universal Darwinism’.

UNIVERSAL DARWINISM

The process of natural selection can be thought of as a simple algorithm: if you have variation, selection, and heredity, then you must get evolution (Chapter 10). This means that evolution can work on anything that is varied and selectively copied. In other words, there can be other replicators and other evolutionary systems. This is the principle of universal Darwinism.

Are there any other evolutionary processes? The answer is yes. Many processes once thought to work by instruction or teaching turn out to work by selection from pre-existing

variation. This is true of the immune system and of many aspects of development and learning (Gazzaniga, 1992). For example, the development of young brains involves the selective death of many neurons and connections; learning to speak involves generating all kinds of strange noises and then selecting from those. Dennett (1995a) provides an evolutionary framework for understanding the various design options available for brains, with each level empowering the organisms to find better and better design moves. He calls it the ‘Tower of Generate-and-Test’. At each level, new variants are generated and then tested. By using the same Darwinian process in new ways, new kinds of minds are created.

Of particular interest here are Darwinian theories of brain function. One example is Gerald Edelman’s (1989) theory of neural Darwinism or neuronal group selection, which forms the basis for Edelman and Tononi’s (2000a) integrated information theory of consciousness (Chapter 5). It depends on three main tenets. ‘Developmental selection’ occurs when the brain is growing and neurons send out branches in many directions, providing enormous variability in connection patterns. These are then pruned, depending on which are most used, to leave long-lasting functional groups. A similar process of ‘experiential selection’ goes on throughout life, with certain synapses within and between groups of locally coupled neurons being strengthened and others weakened, without changes in the anatomy.

Finally, there is the novel process of 're-entry', a dynamic process in which selective events across the brain's various maps can be correlated. Re-entrant circuits entail massively parallel reciprocal connections between different brain areas, allowing diverse sensory and motor events to be synchronised. The activity of groups of neurons can contribute to consciousness if it forms part of what they call the 'dynamic core'. This is an ever-changing yet highly integrated functional group involving large numbers of widely distributed thalamocortical neurons with strong mutual interactions. According to Edelman and Tononi, these principles provide the basis for understanding both the ongoing unity and the endless variety of human consciousness.

This theory certainly entails both selection and variation, but it is not clear that it includes any principle of heredity. In Edelman's theory, variant patterns are generated and selected, but there seems to be no mechanism for copying variants to make new ones. Put another way, there is no replicator. This probably applies also to Crick and Koch's (2003) idea of competing coalitions of neurons. Coalitions vary and compete with each other for dominance, and in that sense are selected, but they are not copied.

William Calvin (1996) is an American neuroscientist whose theory does include such copying. He describes the brain as a Darwin machine and sets even higher standards for a truly Darwinian creative process, listing six requirements, all of which are satisfied in the brain. Most important is his understanding of copying. Throughout the cerebral cortex, he argues, are spatiotemporal firing patterns that represent concepts, words, or images. These patterns depend on the way cortical cells are wired up in columns, and with both lateral inhibition and lateral excitation at different distances. The result is hexagonal structures about half a millimetre across that can be copied or cloned and that compete for survival in a truly Darwinian process. Imagine a vast, ever-shifting quilt of hexagons, all jostling to survive and get copied, some cloning whole areas like themselves while others die out. Consciousness is the currently dominant patch of the hexagonal patterned quilt. This model has not led to further research, but is the kind of principle we would need for a truly Darwinian understanding of the brain.

These theories all deal with Darwinian processes within one brain. Dennett's Tower of Generate-and-Test shows a hierarchy of ways for brains to evolve by reacting to the situations they find themselves in: from creatures hardwired to do only what their phenotypes allow, to creatures that blindly try random options and learn which ones work best, to creatures that use their imaginations to rule out 'truly stupid' options before trying them. Our last theory of universal Darwinism deals with copying between one brain and another, achieved at the top level of Dennett's tower: here culture (from tool use to artistic creations) both requires and enhances intelligence. Creatures at this level can share information and skills with each other, and the first steps have now been taken towards creating them in robot form, using communication between individual robots that have internal models of both themselves and the world (Winfield, 2017).

MEMES AND MINDS

Memes are ideas, skills, habits, stories, or any kind of information that is copied from person to person. They include written and spoken words,

'Talk of memes is just the latest in a succession of ill-judged Darwinian metaphors'

(John Gray, 2008)

rules like driving on the left (or the right), and habits like eating with chopsticks (or a knife and fork), as well as songs, dances, clothes fashions, and technologies. The theory of memes is highly controversial and has been criticised by biologists, sociologists, anthropologists, and philosophers (Aunger, 2000; Richerson and Boyd, 2005; Wimsatt, 2010). Nevertheless, it potentially provides a completely new way of understanding the evolution of consciousness.



MEMES

Origins. Dawkins (1976) coined the term to provide an example of a replicator other than the gene: a cultural replicator.

Definitions

Meme. (mi:m), *n. Biol.* (shortened from *mimeme* [. . .] that which is imitated, after *GENE n.*) 'A cultural element or behavioural trait whose transmission and consequent persistence in a population, although occurring by non-genetic means (esp. imitation), is considered as analogous to the inheritance of a gene' (*Oxford English Dictionary*, January 2018). A meme is any information that is copied from person to person. Many mental events, including perceptions, visual memories, and emotions, are not memes because they are not acquired by imitation or copying. Skills acquired by individual learning, such as avoiding flames or hot chilli, are not memes. Your skateboard is a meme (it was copied), and the idea of skateboarding is a meme, but your skill in riding it is not (you had to learn by trial and error and so does your friend who watches you enviously).

Memeplex. Abbreviated from 'co-adapted meme complex': a group of memes that are passed on together. Memeplexes form whenever a meme can replicate better as part of a group than it can on its own. Memeplexes range from small groups of words, such as sentences and stories, to religions, scientific theories, and works of art, or financial and political institutions.

CONCEPT 11.2

The term 'meme' was coined by Dawkins (1976) to illustrate the principle of universal Darwinism and to provide an example of a replicator other than the gene (Concept 11.2).

Memes count as replicators because they are information that is copied with variation and selection. Of all the thousands of jokes you've ever heard, you have probably remembered very few, and passed on even fewer. For every bestselling book, millions of copies of unpopular ones sit unread on the shelves. As for internet memes – only the funniest versions of Doge's bad grammar get copied millions of times; only the best Gangnam Style dances have been watched billions of times. Memes are copied by imitation, teaching, and reading, and by all the computerised processes of the modern information age. Sometimes they are copied perfectly, but often variation is introduced. This can happen when the copying is imperfect, as in forgetting or misremembering the punchline to a joke, or when old memes are combined in new ways to produce new memes, like all the variations on the 'why did the chicken cross the road?' or 'how many Xs does it take to change a lightbulb?' jokes, or many of the most successful internet memes. This means that the whole of human culture can be seen as a vast new evolutionary process based on memes, and human creativity can be seen as analogous to biological creativity. On this view, biological creatures and human inventions are both designed by the evolutionary algorithm. Human beings are the meme machines that store, copy, and recombine memes (Blackmore, 1999).

The theory of memetics did not start by analogy with genes, although it is often described that way (Searle, 1997; Wimsatt, 2010). Rather, memes are one kind of replicator and genes are another. Analogies can be drawn between them, but often these are not close because the two replicators work quite differently (Blackmore, 2010). For example, genes are based on information stored in molecules of DNA, and copied with extremely high fidelity, while

memes depend on the variable-fidelity copying of human interactions.

Among the similarities are that both genes and memes compete selfishly to be copied, their only interest being self-replication. Some memes succeed because they are useful to us, such as the vast memplexes of technology and the arts and sciences. At the other end of the spectrum are memes that use tricks to get themselves copied. Many of these are essentially 'copy-me' instructions backed up with threats and promises, such as email viruses, pyramid schemes, and religions (Dawkins, 1976). In the middle are vast swathes of culture that are sometimes useful and sometimes destructive, like political and financial institutions. Based on these principles, memetics has been used to explain many aspects of human behaviour and human evolution, including the origins of our big brains and our capacity for language (Blackmore, 1999). A model of creativity involving the evolutionary processes of 'blind variation' (generating ideas in divergent thinking) followed by 'selective retention' (convergent thinking to refine specific ideas) seems to align with brain data derived by several different methods, and involves activity in the default mode network during the first phase (Jung et al., 2013).

The concept of memes is a central part of Dennett's theory of consciousness. He describes a person as 'the radically new kind of entity created when a particular sort of animal is properly furnished by – or infested with – memes' (1995b, p. 341), and a human mind as 'an artifact created when memes restructure a human brain in order to make it a better habitat for memes' (p. 365). On his view, the human brain is a massively parallel structure that is transformed by its infection with memes into one that *seems* to work as a serial machine. Just as you can simulate a parallel computer on a serial one, so the human brain simulates a serial machine on parallel machinery. He calls this the 'Joycean machine' after James Joyce's stream-of-consciousness novels, which tried to convey the parallelism of consciousness through the seriality of language. So, with this virtual machine installed, we come to think about one thing after another, and to use sentences and other mental tools, in a way that suits language-based memes.

This is how the self, the 'centre of narrative gravity' (Chapter 16), comes to be constructed: 'our selves



FIGURE 11.11 • St Paul's cathedral is a meme-spreading monument. The beautiful vistas, awesome dome, inspiring paintings, and delightful music all make people want to worship there, and in the process they spread the memes of Christianity.

Selfplex. A memplex which is formed when people use language that includes references to self. Sentences such as 'I believe x', 'I think y', 'I hate z' give an advantage to memes x, y, and z over simply stating them. In the process, they contribute to belief in an 'I' who has the beliefs, thoughts, and desires. Although the original function was to spread the memes, we now use self-referential language to express many non-memetic ideas, too (e.g. 'I feel angry').

Viral memes. Some memes succeed because they are true or useful or beautiful, while others use tricks to persuade people to copy them. Viral memes include email viruses, Ponzi schemes, and ineffective diets and therapies. Dawkins calls religions 'viruses of the mind' because they infect people by using threats and promises, trick them by discouraging doubt, and reward them for passing on the memplex.

Internet memes. Images, videos, or texts, often humorous or surprising, that are copied, sometimes with deliberate variations, and passed on to potentially millions of others by internet users.

Tremes. Technological memes that are copied, varied, and selected by machines without human involvement.



PRACTICE 11.2

IS THIS A MEME?

As many times as you can, every day, ask yourself '**Am I conscious now?**' Take whatever you were conscious of and ask '**Is this a meme?**' Anything you copied from someone else is a meme, including thoughts in words. Anything that is purely your own and not copied is not. How often is your awareness free of memes?

'Human consciousness is itself a huge complex of memes'

(Dennett, 1991, p. 210)

IS THIS A MEME?

'We, alone on earth, can rebel against the tyranny of the selfish replicators'

(Dawkins, 1976, p. 201)

'there is no one to rebel'

(Blackmore, 1999, p. 246)

have been created out of the interplay of memes exploiting and redirecting the machinery Mother Nature has given us' (Dennett, 1995b, p. 367). The self is a 'benign user illusion of its own virtual machine!' (Dennett, 1991, p. 311).

But perhaps the illusion is not so benign after all. Another possibility is that this illusion of self is actually harmful to *us*, although it benefits the *memes* that make it up. On this view, the self is a powerful memeplex (the selfplex) that propagates and protects the memes within it, but in the process gives rise to the illusion of free will, and to selfishness, fear, disappointment, greed, and many other human failings. Perhaps without it we might be happier and kinder people, although it is hard to imagine consciousness without a self (Chapter 18).

On Dennett's view, 'Human consciousness is *itself* a huge complex of memes (or more exactly, meme-effects in brains)' (1991, p. 210), but this presents two problems. First, memes, by definition, can be copied. Yet our own conscious experiences cannot be passed on to someone else; that is the whole problem and fascination of consciousness. Second, the memes can, arguably, be dropped without consciousness disappearing. For example, at moments of shock, or when silenced by the beauty of nature or in deep meditation, the mind seems to stop. Far from losing consciousness, as Dennett's theory would imply, people say that they become *more* conscious at such moments. This suggests that perhaps human consciousness is distorted into its familiar self-centred form by the memes, rather than that it *is* a complex of memes (Blackmore, 1999). If so, what is left when the memes go away?

Dawkins believes that 'We, alone on earth, can rebel against the tyranny of the selfish replicators' (1976, p. 201), and Csikszentmihalyi urges us to 'achieve control' over our minds, desires and actions: 'If you let them be controlled by genes and memes, you are missing the opportunity to be yourself' (1993, p. 290). But evolutionary processes are not controllable by the creatures they give rise to; and in any case, who is this self who is going to rebel?

Finally, if memes are a second replicator that is copied by the vehicles of the first, could the same thing happen again? Could meme vehicles made by human meme machines, such as computers or phones, become copying machines for a third replicator which we might call *temes* or *tremes*? Certainly the invention of computers, and especially the internet, has brought the concept of memes into popular culture and led to new research in memetics (Shifman, 2013). So, perhaps a new replicator is already evolving in cyberspace, unseen by us, yet supported

by all our interlinked computers and servers that are constantly copying, varying, and selecting vast amounts of digital information (Blackmore, 2010). If this is the right way of looking at the evolution of information technology, we can only speculate whether it might give rise to a new kind of digital consciousness, or perhaps a new illusion of consciousness.



Barlow, H. (1987). The biological role of consciousness./Humphrey, N. (1987). The inner eye of consciousness. In C. Blakemore and S. Greenfield (Eds), *Mindwaves* (pp. 361–374, 377–381). Oxford: Blackwell.

Asking what selective advantage consciousness confers, both authors conclude the advantages are social, but Barlow is more sceptical than Humphrey about introspection.

Blackmore, S. (2010). Memetics does provide a useful way of understanding cultural evolution./Wimsatt, W. (2010). Memetics does not provide a useful way of understanding cultural evolution. In F. Ayala and R. Arp (Eds), *Contemporary debates in philosophy of biology* (pp. 255–272, 273–291). Chichester: Wiley-Blackwell.

Debate on the value of memetics (for and against).

Blackmore, S. (2017). Untestable claims and the evolution of consciousness. *Trends in Ecology and Evolution*, 32(5), 311–312.

A critical review of Feinberg and Mallatt (2016), with a detailed response by the authors at [www.cell.com/trends/ecology-evolution/comments/S0169-5347\(17\)30036-8](http://www.cell.com/trends/ecology-evolution/comments/S0169-5347(17)30036-8).

Graziano, M. A., and Kastner, S. (2011). Human consciousness and its relationship to social neuroscience: A novel hypothesis. *Cognitive Neuroscience*, 2(2), 98–133.


Proposes that consciousness is information created by the neural machinery which evolved for social perception.

Hofstadter, D.R. (2007). A fleeting encounter with zombies and dualism [excerpt]. In D.R. Hofstadter, *I am a strange loop* (pp. 342–349). New York: Basic Books.

Includes sections on ‘Why couldn’t we all just have been zombies?’ (introducing racecar power), ‘Liphosophy’ (belief in a terribly important nonphysical property), ‘Problems with consciousness as a capitalized essence’, and ‘A sliding scale of élan mental’.

Macphail, E.M. (2009). Evolution of consciousness. In T. Bayne, A. Cleeremans, and P. Wilken (Eds), *The Oxford companion to consciousness* (pp. 276–279). Oxford: Oxford University Press.

Uses the experience of pain as a case study in the difficulties of making sense of consciousness and evolution.



The evolution of machines

CHAPTER

TWELVE

MINDS AND MACHINES

Is there something special about human beings that enables us to think, see, hear, feel, and fall in love; that gives us a desire to be good, a love of beauty, and a longing for something beyond? Or are all these capacities just the products of a complicated mechanism? In other words, am I just a machine?

Is there something about human-made machines that means they can never think, feel, intuit, fall in love, and long for something beyond? Or could they one day do all those things and more? In other words, could there be machine consciousness (MC) or artificial consciousness (AC)?

If there could, we may have some kind of moral responsibility for our creations. We may also find that their existence changes our views of our own consciousness.

'work on machine consciousness may come strongly to affect how consciousness is seen'

(Clowes, Torrance, and Chrisley, 2007, p. 14)

MINDS LIKE MACHINES

The suspicion that we humans are really machines has a long history. The early Greek materialists Leucippus, Democritus, and Lucretius, for example, argued that there exist only atoms and the void, and everything happens through natural processes. This mechanical view seemed to exclude divine creation and threaten free will, and was rejected by Plato and Aristotle, who argued for worlds and forces beyond the material world.

In the seventeenth century, Descartes argued that the human body was a mechanism but that no mechanism alone was capable of speech and rational

• SECTION FOUR : EVOLUTION

thought – for that, *res cogitans* or thinking-stuff was needed (Chapter 1). Among those who rejected his dualism was Gottfried von Leibniz, best known for his work on calculus and his philosophy that all matter consists of simple nonmaterial substances, which he called little minds, or monads. This meant that he rejected materialism, and he justified this with his famous allegory of the mill (1714/1965). Imagine a machine whose construction enabled it to think, feel, and perceive. Imagine, then, that the machine were enlarged while retaining the same proportions, so that we could go inside it, like entering a windmill. Inside we would find only pieces working upon one another and never anything to explain the perception. From this he concluded that to explain perception, we must look to a simple substance rather than to the workings of a machine, which can never have the unity that consciousness does.

Leibniz's thought experiment can be applied directly to the human brain. Imagine making the neurons bigger and bigger so that we could go inside. What would we see but synapses and chemicals working upon one another? Leibniz also argued that the 'I' could not be found in a mill. With this thought experiment he faced, long before neurons and synapses had ever been heard of, the same questions we face now. How can a machine feel as though it is, or has, a conscious self?

Another of Descartes's critics took the opposite tack and scandalised the world with his infamous book *L'Homme Machine* (*Machine Man, or Man-Machine*, 1748). Julien Offray de la Mettrie was a pleasure-loving French philosopher and physician who rejected Descartes's break between man and the soulless animals, and classified humans as living machines. His materialist and irreligious views provoked outcry, especially since they led him to a morality based on rejecting guilt and seeking pleasure, and he was forced to flee France, first to the Netherlands and then to Berlin.

The idea that we are machines has never seemed comfortable; but now that we understand so much more biology and psychology, the question is not so much 'Am I a machine?' but 'What kind of machine am I?', and, for our purposes here, 'Where do "I" fit in?' and 'Where does consciousness fit in?'

There are two ways to seek answers. We can start with the biology and try to understand how natural systems work, or we can build artificial systems and see how far they can match human abilities. As Stevan Harnad (2007) describes it, we can reverse-engineer the brain to see how it works, or we can forward-engineer a brain by building something that can do what brains do.

In consciousness studies, the two endeavours are converging. From the natural direction, science has successively explained more and more of the mechanisms of perception, learning, memory, and thinking, and in so doing has only amplified the ancient open question about consciousness. That is, when all these abilities have been fully explained, will consciousness be accounted for too, or still be left out?

From the artificial direction, better and better machines have been developed, leading to the obvious question of whether they are conscious already, or could be one day. If machines could do all the things we do, just as well as we do them, would they be conscious? How could we tell? Would they *really* be conscious, or just zombies simulating consciousness? Would they *really* understand what they

'It's the way our machine is provisioned that makes us lively or brave'

(de la Mettrie, 1748/2015, trans. Bennett, p. 5)

said and read and did, or would they just be acting *as if* they understood? We arrive at the same question – is there something extra that is left out?

These are some of the central questions for this chapter. While the main objective is to think about AC, this has been so closely bound up with the topic of artificial intelligence (AI) that we need to begin there. To us, adding machines and other automata may not seem to have anything to do with consciousness, but we should appreciate that rationality was long prized above all other qualities of the human mind and assumed to be a product, perhaps the highest product, of human consciousness. In fact it turned out that rational, logical thinking is far easier for artificial machines than are some of the things that animals do easily, like seeing, finding food or mates, and showing emotions. So we no longer assume that rationality is a sign of consciousness, and may be less impressed by mathematical machines even though it is with them that humans began to think about and create artificial consciousness.

MIND-LIKE MACHINES

From the fourth century BC the Greeks made elaborate marionettes, and later complete automatic theatres, with moving birds, insects, and people, all worked by strings and falling-weight motors. These machines mimicked living things in the sense that they moved like them, but it was not until much later that the idea of thinking machines became possible.

In 1642, the French philosopher and mathematician Blaise Pascal began work on one of the first ever digital calculating machines when he was only 19 years old. Although it could add and (with difficulty) subtract, using interconnected rotating cylinders, it was too cumbersome to be commercially useful. In 1672 Leibniz developed a machine that could add, subtract, multiply, and divide, although it too was unreliable. Commercially successful machines did not appear until the nineteenth century.

During the eighteenth century, automata became immensely popular, with the most famous including a flute-playing boy, a duck with a digestive system, and the earliest chess-playing machine, the ‘Turk’. This consisted of a wooden cabinet with doors that opened to show cogs and wheels inside, and an impressive life-size wooden figure that wore a robe and turban and used mechanical hands to move chess pieces on a board. The Turk was said to beat most challengers within half an hour and toured the great cities of Europe for decades without its trick being exposed. But a trick it certainly was (Standage, 2002).

Automata continued to fascinate and frighten, and in 1818 Mary Shelley captured this fear in her novel about Frankenstein and his gruesome monster. But soon the technology began to be used for more scientific purposes.

In the 1830s, the English mathematician Charles Babbage was infuriated by unreliable mathematical tables and conceived the idea of a ‘difference engine’ that could compute the tables accurately and even print them. It was never completed, and the even more ambitious ‘Analytical engine’ was never even started. This was to have had a processing unit of cogs and wheels controlled by punched cards, like those used in looms for weaving cloth, which would have allowed it to carry out many different functions. This was probably not technically feasible

‘Every intelligent ghost must contain a machine – an information-processing machine’

(Slovan, 2014, p. 1)




FIGURE 12.1 • The mechanical Turk was the first ever chess-playing machine. His hands were moved by intricate machinery under the table, but the real player was hidden inside. The online crowdsourcing marketplace is named after him (Photo: AKG London).

at the time, yet it has taken its place in history as the first general-purpose, programmable calculating machine.

Among the ideas that were fundamental to such machines was Boolean algebra, invented by the English mathematician George Boole. As a young man, working as an assistant teacher in Doncaster in 1833, Boole went walking one day on the Town Fields. There he had a sudden insight: one of the famous 'eureka' moments of science. He saw that just as mathematical principles could explain the function of cogs in machines, so they might be able to explain what he called 'the laws of thought', and he believed that in this way mathematics might solve the mysteries of the human mind. He showed how logical problems could be expressed as algebraic equations and therefore solved by mechanical manipu-

lation of symbols according to formal rules. This required only two values, 0 and 1, or false and true, and the rules for combining them. Boole did not succeed in solving the mystery of mind, as he had hoped, but Boolean algebra was fundamental to the computer revolution.

In the 1930s, the American mathematician and founder of information theory, Claude Shannon, realised that Boolean algebra could be used to describe the behaviour of arrays of switches, each of which has only two states, on or off. He used a binary code and called each unit of information a 'binary digit' or 'bit'. All this made possible the idea that logical operations could be embodied in the workings of a machine.

As so often happens, it was the pressures of war that drove on the invention of computing machinery. The first general-purpose computers were built in the Second World War to decode German ciphers and to calculate the tables needed to guide ballistics. The master code-breaker, though this was only revealed thirty years after the war ended, was the brilliant English mathematician Alan Turing.

Turing worked on algorithms, that is, sets of step-by-step instructions for operations to be performed. Problems are said to be 'computable' if they can be formulated and solved by using an appropriate algorithm. Turing proposed the idea of a simple machine that could move an indefinitely long tape backwards and forwards one square and print or erase numbers on it. He showed that this simple machine could specify the steps needed to implement *all* computational algorithms.

The principle underlying this is an abstract machine, now known as the Turing machine. An important aspect of this idea is that the abstract machine has 'multiple realisability' and 'substrate-neutrality'. That is, it can use tapes or

chips, or be made of brain cells, beer cans, water pipes, or anything else at all, as long as it carries out the same operations. This gives rise to the idea of the Universal Turing Machine, a machine that can, in principle, imitate any other Turing machine. The 'in principle' is needed because the machine may require an unlimited memory store and unlimited time in which to do its calculations. Even so, modern computers can be thought of as universal Turing machines, since many different 'virtual machines', such as word processors, web browsers, or spreadsheets, can be run on the same physical machine; even PowerPoint has been shown to be able to simulate any Turing machine.

Even the slow and cumbersome early computers inspired comparisons with the human mind. During the Second World War, the Cambridge psychologist Kenneth Craik began to develop the idea that the human mind translates aspects of the external world into internal representations and that perception, thought, and other mental processes consist of manipulating these representations according to definite rules, as a machine might do. He died in a car crash at the age of 31, but these ideas became one of the dominant paradigms in psychology for the rest of the century, giving rise to the idea that what we are conscious of is these internal representations or mental models – in other words, that the contents of consciousness are mental representations.

Although computers rapidly became faster, smaller, and more flexible, initial attempts to create AI depended on a human programmer writing programs that told the machine what to do using algorithms that processed information according to explicitly encoded rules. This is now referred to – usually by its critics – as GOFAI (pronounced 'goofy'), or 'Good Old-Fashioned AI'.

One problem for GOFAI is that human users treat the processed information as symbolising things in the world, but these symbols are not grounded in the real world for the computer itself. So for example, a computer might calculate the stresses and strains on a bridge, but it would not know or care anything about bridges; it might just as well be computing stock-market fluctuations or the spread of a deadly virus. Similarly, it might print out plausible replies to typed questions without having a clue about what it was doing. Because such machines merely manipulate symbols according to formal rules, this traditional approach is also called rule-and-symbol AI.



PROFILE 12.1

Alan Turing (1912–1954)



Born in London and educated at Cambridge, Alan Turing was an extraordinarily brilliant mathematician. He is often called the father of both computer science and artificial intelligence, partly because of his famous work on computable numbers which

led to the idea of the Universal Turing Machine. He also created the Turing Test, which pits a machine against a person as a way of finding out whether the machine can think. Thirty years after the Second World War, Turing was revealed as the master code-breaker who had broken the famous Enigma cipher. He also created the first functioning programmed computer, the Colossus, to read the highest-level German secret codes. He was homosexual, and was eventually arrested and tried for what was then illegal behaviour, and forced to take female hormones. He died in June 1954 of cyanide poisoning, probably by suicide. He was granted a posthumous Royal pardon in 2013.

From this emerged the 'computational theory of mind'. As Searle later put it:

Many people still think that the brain is a digital computer and that the conscious mind is a computer program, though mercifully this view is much less wide-spread than it was a decade ago. Construed in this way, the mind is to the brain as software is to hardware.

(1997, p. 9)

Searle distinguished two versions of this theory: 'Strong AI' and 'Weak AI'. According to Strong AI, a computer running the right program would be intelligent and have a mind just as we do. There is nothing more to having a mind than running the right program. Searle claimed to refute this with his famous Chinese Room thought experiment (which we look at later in this chapter). According to Weak AI, computers can *simulate* the mind and simulate thinking, deciding, and so on, but they can never create *real* mind, *real* intentionality, *real* intelligence, or *real* consciousness, only *as-if* consciousness. This is like a meteorologist's computer that may simulate storms and blizzards but will never start blowing out heaps of fluffy cold snow.

A similar distinction is made between 'Weak AC' (or MC) and 'Strong AC' (or MC). One strand of research uses computational, robotic, and other artificial means to model consciousness, hoping to understand it better: this is Weak AC, Weak MC, or MMC (Machine Modelling of Consciousness; Clowes, Torrance, and Chrisley, 2007). 'The key intention of the MMC paradigm is to clarify through synthesis [of notions from psychology, neuroscience, philosophy, and introspection] the notion of what it is to be conscious' (Aleksander, 2007, p. 89). The other strand, Strong AC, aims to actually construct a conscious machine for its own sake. By analogy with the arguments over AI, we might say that someone who believes in Weak AC thinks we can learn about consciousness by building machines; someone who believes in Strong AC thinks we can create consciousness by building machines.

DEVELOPMENTS IN COMPUTING

According to Moore's Law on Integrated Circuits, the number of transistors on a chip doubles every two years. Remarkably, this observation (not really a true



CONCEPT 12.1

BRAINS AND COMPUTERS COMPARED

Digital v. analogue. The vast majority of computers are digital, even though they may simulate analogue processes. A digital system works on discrete states, whereas an analogue system works on continuous variables. In music recording, for example, digital CDs code the frequency and intensity of sound (a naturally analogue signal) by discrete digits, whereas analogue vinyl records represent them by contours in the groove. Digital coding makes higher-fidelity copying possible because slight variations are automatically eliminated as long as they are not large enough to switch a 0 to a 1, or vice versa.

Is the human brain digital or analogue? The answer is both. A neuron either fires (a wave of depolarisation runs along its membrane) or not, and to this extent is digital, yet the rate of firing is a continuous variable. Another analogue process is spatial summation. Imagine an axon with a synapse on a second cell's dendrite (see Figure 12.2). When the first cell fires, neurotransmitters cross the synapse and change the state of polarisation of the post-synaptic membrane briefly and for a short distance around the synapse.

law), made in 1965, seems to have held ever since, though in the past few years the rate of change has slowed to a doubling only every two-and-a-half years, and some people predict saturation within the next decade – though then again, if you apply the prediction to neural networks, it keeps being outstripped. This remarkable expansion describes brute computing power, but there have been more fundamental changes in AI that are relevant to understanding consciousness.

CONNECTIONISM

The 1980s saw the flowering of ‘connectionism’, a new approach based on artificial neural networks (ANNs) and parallel distributed processing. Part of the motivation was to model the human brain more closely, although even twenty-first century ANNs are extremely simple compared with human brain cells. The many types of network include recurrent, associative, multilayered, and self-organising. The big difference from GOFAI is that ANNs are not programmed: they are trained. To take a simple example, imagine looking at photographs of people and deciding whether they are male or female. Humans can do this easily (although not with 100% accuracy) but cannot explain how they do it. So we cannot use introspection to teach a machine what to do. With an ANN, we don’t need to. In supervised learning, the system can be shown a series of photographs and for each one produce an output: male or female. If this is wrong, the synaptic weights are adjusted and the network is shown the next, and so on. Although it begins by making random responses, a trained network can correctly discriminate new faces, as well as ones it has seen before.

How does it do this? Even a simple network consists of many units, each resembling a neuron in the sense that it sums the inputs it receives according to a mathematical function, and produces an output (a ‘1’ or a ‘0’, fire or don’t). The units are connected in parallel, each connection having a weight, or strength, that can be varied. A simple network might consist of three layers: an input layer, a hidden layer, and an output layer. For the example of faces, the input layer would need enough units to encode an array corresponding to the photographs (e.g. one for each pixel), and the output layer would need one unit that outputs ‘0’ for male, ‘1’ for female. For




FIGURE 12.2 • The state of polarisation of any part of the postsynaptic membrane varies continuously depending on the effects of many synapses at different distances (an analogue process). When the polarisation at the cell body reaches a critical threshold, the second cell fires (a digital process).

Now imagine the effects of lots of other synapses on the same cell but at slightly different times and distances from the cell body. These all add up so that if the polarisation at the cell body reaches a critical threshold, the second cell fires. The process of summation is analogue, but the final output – to fire or not to fire – is digital. It is not possible to characterise the brain as simply either digital or analogue.

Serial v. parallel. Many digital computers, and certainly all the early ones, process information very fast, but serially, i.e. one thing after another. They have a single central processing unit and can work simultaneously on different tasks only by dividing the tasks up and switching between them. By doing this a serial machine can simulate a parallel machine.

By comparison, neurons operate very slowly, but the brain is massively parallel, with no central processor but millions of simultaneously operating cells. To some extent, this massive parallelism compensates for lack of speed. Even so, the overall computing speed of the latest supercomputers is estimated at four times as fast as a human brain.

Brains have different areas for vision, hearing, planning, and so on, all operating in parallel all the time, and within small areas of the brain patterns of information move about in complex networks that have no serial organisation. The brain does seem to have bottlenecks, though, such as limited short-term memory and attention (Chapter 7). Also many outputs, including spoken and written language, are serial. In this sense, the brain is a parallel machine simulating a serial machine; this is Dennett’s Joycean machine (Chapter 11).

Computable v. non-computable.

A computable procedure is one that can be described explicitly, and any such procedure can be carried out by a computer program (this is the Church–Turing thesis). Computational functionalism is the doctrine that the brain is essentially a Turing machine and its operations are computations. If this is true, then it should be possible to reproduce all human abilities by carrying out the right computations, making strong AI feasible. Against this is the claim that such computations would only *simulate* human functions; that there is more to consciousness than running the right program. Turing himself showed that some functions are noncomputable and Penrose argues that mathematicians can intuitively see noncomputable truths, meaning that the brain is not a Turing machine and conscious understanding goes beyond computation. (See the website for more on the maths of noncomputability.)

Deterministic v. nondeterministic.

A machine that always produces the same outcome from the same input and the same internal state is deterministic; one that can produce different outcomes is nondeterministic. Digital computers are deterministic. Note that this does not mean that their outcome must be predictable. For example, chaos theory shows that for some deterministic processes, the outcome varies dramatically with only very slight differences in initial conditions. Nor does it mean that computers cannot be creative. The evolutionary algorithm (Chapter 10) is *par excellence* a deterministic procedure that yields creativity. Computers can simulate nondeterministic systems by adding pseudo-randomness.

Brains, at least at one level, are nondeterministic. They are warm, wet, and noisy, and therefore cannot always produce the same output to the same input. Neurons are living entities whose electrical properties change as their dendrites grow or move. Synapses form and decay, and their strength changes in response to use. So the machine itself is never the same from one moment to the next. At a smaller scale, though, the underlying molecular processes are usually assumed to be deterministic. This is one reason why there appears to be no room for free will, and adding randomness, as one can do with a computer, does not provide a meaningful kind of ‘freedom’ (Chapter 9). Going smaller still, one reaches the level of quantum effects and quantum indeterminacy. Some have argued that this is the ultimate source of human creativity, free will, and consciousness.

a more complex task, such as identifying individual faces, it would need enough output units to encode any allowable identity. During training, a program compares the net’s actual output with the correct output and makes adjustments to the weights accordingly – but how? The best-known method uses the back-propagation algorithm (meaning that the error is iteratively fed back into the network to update the weights). As training proceeds, the errors get gradually smaller until the network responds more or less correctly. If the training set of photographs is appropriately chosen, the network should now perform well on a completely new photograph.

Note that the process of adjusting the weights is algorithmic, or rule-based, and the whole system may be run on a digital computer, or hardcoded on to a chip for far greater speed. The system contains nothing that tells it how to recognise men and women. The ANN works this out for itself, and even its creators cannot know what exactly the weights mean. Unlike traditional machines, connectionist networks do not just do what their programmers tell them to do. This is a long way from good old-fashioned rule-and-symbol AI, and getting further away all the time, as new developments bring in fuzzy logic (allowing ANNs to take into account concepts like ‘usually’, ‘somewhat’, and ‘sometimes’ rather than just binary true/false values) and explore the possibility of pulsed neural networks to mimic how biological neural networks use the timing of pulses to communicate information and perform computations. Deep learning (in networks with many layers) has also been accelerated by the advent of massively parallel graphics processors (GPUs) developed for video gaming. These are used to drive applications that require vast processing power to train billions of ‘software-neurons’.

ANNs are useful for many purposes, including recognising handwriting, controlling robots, mining data, forecasting market fluctuations, and filtering spam, and may soon be used in many more applications like self-driving cars. The connectionist–computational debates continue, but so does the gradual movement from understanding cognition as manipulation of static symbols towards treating it as a continuous dynamical system that cannot be easily broken down into discrete states.

EMBODIED COGNITION

The machines described so far are all disembodied, confined inside boxes and interacting with the world only through humans. When first put to work controlling robots, most could carry out only a few simple, well-specified tasks in highly controlled environments, such as in special block worlds in which they had to avoid or move the blocks. This approach seemed sensible at the time because it was based on an implicit model of mind that was similarly disembodied. It assumed the need for accurate representations of the world, manipulated by rules, without the messiness of arms, legs, and real physical problems. We might contrast this with a child learning to walk. She is not taught the rules of walking; she just gets up, falls over, tries again, bumps into the coffee table, and eventually walks. By the same token, a child learning to talk is not taught the rules either; in the early days she pieces together fragments of sounds she hears and gestures she sees, parses words wrong, and eventually makes herself understood.

The connectionist approach is far more realistic than GOFAI, but still leaves out something important. Perhaps it matters that the child has wobbly legs, that the ground is not flat, and that there are real obstacles in the way; maybe it matters that she has the vocal cords she does and that her parents' gestures are constrained by their limbs.

As we saw in Chapters 5 and 6, embodied or enactive or 4E cognition are names for the general idea that mind can be created only by interacting in real time with a real environment – the idea, drawing on the phenomenology of Merleau-Ponty, 'that cognition is not the representation of a pre-given world but is rather the enactment of a world and a mind' (Varela, Thompson, and Rosch, 1991, p. 9). Andy Clark (1997) wants to put brain, body, and world together again – both causally and computationally speaking. 'Fortunately for us,' he says, 'human minds are not old-fashioned CPUs trapped in immutable and increasingly feeble corporeal shells. Instead, they are the surprisingly plastic minds of *profoundly* embodied agents' (2008, p. 43). What he means by '*profoundly* embodied' is that every aspect of our mental functioning depends on our intimate connection with the world we live in. Our 'supersized' minds and our powers of perception, learning, imagination, thinking, and language are all created by brains interacting with bodies and their environments, both physical and social.

On this view the real world is far from being a messy complication we can do without; rather, it provides the very constraints and feedback that make perception, intelligence, and consciousness possible. Human intelligence is not just 'recognition intelligence': it is about using understanding to make autonomous real-time decisions. Creating machines this way means constructing real, physical, autonomous agents that move about in the real messy world, working from the bottom up rather than the top down. There is no point in a driverless car recognising a collection of pixels as a white van slowing down quickly unless it




FIGURE 12.3 • This artificial neural network (ANN) has just three layers of units: the input layer, the output layer, and a hidden layer in between. During training, the weights on the connections between the units are adjusted until the network provides the correct output. Such a network can learn to recognise faces, to produce sounds in response to written text, and many other tasks, depending on what is connected to the input and output units.

'human minds are not old-fashioned CPUs trapped in immutable and increasingly feeble corporeal shells'

(Clark, 2008, p. 43)



FIGURE 12.4 • A termite mound in West Bengal, India. Each individual termite follows simple rules about when to add mud and when to remove it. None has a plan of the overall mound, yet the complex system of tunnels and walls emerges. Is consciousness an emergent phenomenon like this?

can assess the current situation and take evasive action. This approach is sometimes called situated robotics, or behaviour-based (as opposed to knowledge-based) robotics.

One implication is that intelligent behaviours can emerge from simple systems, perhaps holding out the hope that consciousness might do the same. There are many examples of such emergence in biology. For example, termites build extraordinary structures that look as though they must be planned, when in fact they emerge from simple rules about when to add or remove mud, embodied in the individual termites. Emergent intelligence in social insects is the inspiration behind the field of swarm robotics (Brambilla et al., 2013), in which large numbers of simple robots following relatively simple rules can produce multiple complex swarm behaviours,

whether for use in medicine, disaster rescue, or autonomous warfare.

As for single robots – imagine watching a small wheeled robot moving along next to a wall. It does not bump the wall or wander far away from it, but just wiggles along, reliably following the bends and turning the corners. How? It might have been programmed to follow the wall using a detailed internal representation of the area and instructions for coping with each eventuality, but in fact it need not be. All it needs is a tendency to veer to the right, and a sensor on the right side to detect close objects and make it turn slightly to the left whenever it does so. By balancing the two tendencies, wall-following behaviour emerges.

This is a good example of that slippery concept, an emergent property. An apparently intelligent behaviour has emerged from an extremely simple system. This might help us to consider whether consciousness could also be an emergent property of a physical system, as some believe it is (Humphrey, 1987; Mithen, 1996; Searle, 1997; Feinberg and Mallatt, 2016).

INTELLIGENCE WITHOUT REPRESENTATION

Traditional AI assumed that intelligence is all about manipulating representations, yet our wall-following robot managed with none. How much further could this go? To find out, Rodney Brooks and his colleagues at MIT spent many years building robots with no internal representations (Brooks, 1997, 2002).

Brooks's 'creatures' can wander about in complex environments such as offices or labs full of people and carry out tasks such as collecting rubbish. They have several control layers, each carrying out a simple task in response to the environment. These are built on top of each other as needed and have limited connections enabling one layer to suppress or inhibit another. This is referred to as 'subsumption architecture' because one layer can subsume the activity of another. Brooks's robot Allen, for example, had three layers: the lowest prevented him from touching other objects by making him run away from obstructions but

otherwise sit still; the second let him wander around without crashing into things; and the third made him explore by looking for distant places and trying to reach them. Correction signals operated between all three. Such a creature's overall behaviour looks intelligent to an observer but, says Brooks, 'It is only the observer of the Creature who imputes a central representation or central control. The Creature itself has none; it is a collection of competing behaviors' (1997, p. 406).

This is related to Marvin Minsky's (1986) idea of 'the society of mind', in which intelligence emerges from many separate modules doing many simple things all at once; to Ornstein's (1991) description of the mind as a 'squadron of simpletons'; to Dennett's (1991) replacement of the inner audience and 'Central Meander' with a pandemonium of stupid machine-like homunculi; and to Clark's (2013) argument that the mind is best understood as a distributed 'prediction machine'. By building robots this way, Brooks discovered that 'When we examine very simple level intelligence we find that explicit representations and models of the world simply get in the way. It turns out to be better to let the world itself serve as its own model' (1991, p. 396). Although Brooks makes no claims to biological significance, this is the same conclusion that Kevin O'Regan, Alva Noë, and others came to from studying change blindness in humans (Chapter 3). It seems that representations of the world may not always be necessary for building effective robots, and evolution may not have used them when building our vision system, either. Representations are still crucial in other respects: for the sensorimotor theory in storing knowledge about the laws of sensorimotor contingency, and for predictive coding in providing generative models based on prior experience. But the representations are not 1:1 mappings of the 'outside' world.

All this is highly relevant to understanding consciousness. Along with GOFAI goes the idea that conscious experiences are mental models or inner representations of the world. Although intuitively plausible, this idea is problematic. For example, it is not clear how a mental model can *be* an experience, nor why some mental models are conscious while most are not. These are the familiar problems of subjectivity, and of the magic difference between conscious and unconscious processes.

Doing away with representations may solve some problems, but it raises others. In particular, the nonrepresentational approach has difficulties dealing with experiences that are not driven by continuous interaction with the outside world, such as reasoning, imagining, and dreaming. On representational theories, it is easy to think that when I dream of drowning in huge waves, my brain is constructing representations of sea, water, and waves, and simulating death; but if there are no



FIGURE 12.5 • William Grey Walter with one of his famous 'tortoise' robots photographed in 1951. He built two prototypes, Elmer and Elsie, in Bristol in 1948–9. Later six more were built, and displayed at the Festival of Britain in 1951. They had a photocell eye, two vacuum tube amplifiers that drove relays to control steering and drive motors, and a Perspex shell with a switch that operated when the shell contacted anything. They moved about autonomously in a lifelike manner, demonstrating the beginnings of artificial intelligence, and showed a form of self-preserving behaviour by crawling back into their charging hutch when their batteries ran low. In 1995 what was thought to be the last remaining Grey Walter tortoise was found and repaired by Owen Holland, and finally ended up in the Science Museum.

'It turns out to be better to let the world itself serve as its own model.'

(Brooks, 1997, p. 396)

'We are not cognitive couch potatoes idly awaiting the next "input", so much as proactive predictavores'

(Clark, 2015, p. 52)

representations, what could it be doing? This is a challenge for embodied cognition, and for enactive and sensorimotor theories, but there is growing evidence that embodied, enactive, and extended processes contribute to all these activities: that performing congruent actions helps us understand action-based words and concepts (even in 'dead' metaphors like *grasp an idea*), that we move our eyes in similar ways when we imagine and when we see, that we incorporate bodily stimulation into our dreams, and that we use dreaming as a chance to rehearse and optimise the interactive hypothesis-testing of our probabilistic minds.

Interestingly, however, some of the very people who were most attracted to nonrepresentational robotics have discovered that rather than giving robots representations from the outset, they can be built to construct their own internal models. In one example, a wall-following robot builds concepts about itself and the walls it follows to construct a map of its environment, as well as a model of itself, both of which it uses to make decisions about its behaviours by estimating their outcomes (Holland and Goodman, 2003). But robotics researchers Owen Holland and Rod Goodman also point out that as robots become more and more complicated, the challenge of knowing whether an internal model is present, what it corresponds to, and how it is being used gets greater and greater.

There have been many other important developments in AI, but the few covered here at least provide a sketchy outline to guide us when we ask whether a machine can be conscious. This is bound to be a tricky question. How can we know? How can we tell whether we've succeeded? We may get a little help from Turing's famous test for whether a machine can think.



PRACTICE 12.1

AM I A MACHINE?

As many times as you can, every day, ask yourself: **'Am I a machine?'**

The idea is to watch your own actions and consider them in light of the ideas presented here. Are you like a simple autonomous robot? Could a machine created by humans ever do what you are doing now? If so, would the machine feel like you do? You may discover that asking these questions while going about your ordinary life makes you feel more machine-like. What is going on here?

If you find an inner voice protesting 'But I am not a machine!', investigate who or what is rebelling.

THE TURING TEST

Turing's classic paper of 1950 begins, 'I propose to consider the question "Can machines think?"' (p. 433). He dismissed the idea of answering it by analysing the terms 'machine' and 'think' as no better than collecting a Gallup poll of opinions, and proposed instead to base a test on what machines can do.

What, then, is a good test of what a machine can do? Among all the possible tests one can think of, two come up again and again. The first is playing chess. Surely, people have long thought, if a machine can play chess, then it must be intelligent, rational, and able to think. Descartes would presumably have been impressed by such a machine since, like his contemporaries, he prized human rationality far above things that 'lower' animals can do easily, such as walk about and see where they're going. So it is perhaps not surprising that in the early days of computing, it seemed a great challenge to build a computer that could play chess.

After the trick games played by the mechanical Turk, the first serious game took place in Manchester in 1952, with Turing playing the part of a machine against a human opponent. He had written a program on numerous sheets of paper and consulted them at every move, but was easily defeated. In 1958, the first game with an actual machine was played with an IBM computer, and from then on computer chess improved rapidly. Most chess programs relied on analysing several moves ahead. This quickly produces a combinatorial explosion (also known as 'the curse of dimensionality'), because for every possible next move, there are even more possible next moves after that. Programmers and mathematicians invented many ways to get around this, but to some extent computer chess got better just by brute-force computing power. In 1989, the computer Deep Thought took on the world chess champion Gary Kasparov, who told reporters that he was defending the human race against the machine. This time the machine lost, but eventually, in 1997, its successor, Deep Blue, beat Kasparov (for the personal story, see Kasparov, 2017).

Deep Blue consisted of thirty-two IBM supercomputers connected together and could evaluate 100 million positions per second, but no human being plays chess like this. So, was Deep Blue intelligent? Could it think? Searle (1999) said not, arguing that Deep Blue, like the Turk, relied on an illusion, and the real competition was between Kasparov and the team of engineers and programmers. The team said that they never thought their machine was truly intelligent. It was an excellent problem-solver in one domain but could not teach itself or learn from its own games.

In subsequent human-computer battles, another world champion, Vladimir Kramnik, was defeated by Deep Fritz, and a whole team of computers beat a strong human team. More recently, a chess program called Giraffe has been developed not to rely on brute processing power to search through all possible moves, but to evaluate the possibilities and narrow down the avenues worth pursuing, using a deeply layered neural network. Giraffe plays against itself with the aim of improving its ability to predict how it would evaluate a future position relative to the outcomes of winning, losing, or drawing. In this way, Giraffe can bootstrap its way to matching the best chess engines in the world, trained over years often by human grandmasters.

One of the latest developments is to switch chess for the ancient Chinese game of Go, which has a total number of possible moves orders of magnitude larger than the number of atomic particles in the observable universe. In 2016, Google's program called AlphaGo shocked one of the most experienced human players with a move a human would simply never have thought of doing (Wong and Sonnad, 2016), earning itself an honorary 9-dan black belt. The principles AlphaGo was trained on were standard ANN with reinforcement learning, using a prior stage to



FIGURE 12.6 • Progress in artificial intelligence has been dramatic, from Pascal's earliest calculating machine to AlphaGo, shown here playing Lee Se-Do. Does AlphaGo's apparent creativity suggest artificial consciousness?

learn from human players a function that allows it to capture some 'intuitive' sense of good board position.

So can these machines think? They certainly have limitations as well as strengths: they often require lots of ancillary information and large numbers of human examples to learn from, and tiny perturbations can result in crucial misclassifications (Szegedy et al., 2014). But the answer really depends on what you mean by 'think', and investigating that, Turing argued, is not a useful way forward.

Instead, Turing chose something altogether different for his test: whether a computer could hold a conversation with a human. Descartes had claimed this to be impossible and, interestingly, it was one of the tricks attempted by the Turk. In its earliest version, having finished the chess, the Turk would invite people to ask questions and reply to them by pointing at letters on a board. But this was soon dropped from the show. Although audiences could just about believe in a chess-playing automaton, when it claimed to be able to answer questions they assumed it was just a trick and the fascination was lost (Standage, 2002). Perhaps holding a conversation has always been implicitly perceived as harder than playing chess.

The Turk looked like a human, but Turing did not want appearance to confuse his test for a thinking machine, so he cleverly avoided this problem. First he described 'the imitation game', which was already a popular parlour game. The object of this game is for an interrogator or judge (C) to decide which of two people is a woman. The man (A) and the woman (B) are in another room so that C cannot see them or hear their voices and can only communicate by asking questions and receiving typed replies. A and B both try to reply as a woman would, so C's skill lies in asking the right questions.

Turing goes on:

We now ask the question, 'What will happen when a machine takes the part of A in this game?' Will the interrogator decide wrongly as often when the game is played like this as he does when the game is played between a man and a woman? These questions replace our original, 'Can machines think?'

(1950, p. 434)

'at the end of the century the use of words and general educated opinion will have altered so much that one will be able to speak of machines thinking without expecting to be contradicted'

(Turing, 1950, p. 442)

Turing provides a critique of his own test. He points out that it neatly separates the intellectual and physical capacities of a person and prevents beauty or strength from winning the day. On the other hand, it may be too heavily weighted against the machine. A human pretending to be a machine would always fail if tested on arithmetic, and he wonders whether this test is unfair because machines might be able to do something that ought to be described as thinking but that is very different from what a human does. He concludes, though, that if any machine

could play the game satisfactorily, we need not be troubled by this objection. He gives sample questions and answers, and interestingly, these include a chess question, showing how broad and flexible his test is.

Finally, Turing considers many possible objections to the idea that a machine could ever truly be said to think, and states his own opinion on the matter.

I believe that in about fifty years' time it will be possible to programme computers, with a storage capacity of about 10^9 , to make them play the imitation game so well that an average interrogator will not have more than 70 per cent. chance of making the right identification after five minutes of questioning. [. . .] at the end of the century the use of words and general educated opinion will have altered so much that one will be able to speak of machines thinking without expecting to be contradicted.

(1950, p. 442)

What a prescient and carefully worded prediction! Turing was absolutely right about the change in the use of words. I do not expect to be contradicted if I say that my laptop is 'thinking about it' when it reacts slowly, or 'my phone thinks it's Thursday' when the date is wrong. We talk to Siri and our 'OK Google' apps in the expectation that they'll hear, understand, and respond to us. On the other hand, even my lowly desktop has a storage capacity far larger than Turing guessed, and yet it could not pass his test.

When the fifty years were up, many programs could pass limited Turing tests. The first was ELIZA, which used scripts based on rudimentary pattern-matching – repeating and lightly transforming sentences in a psychotherapeutic manner – to give an illusion of understanding, and was genuinely helpful to people with psychological difficulties (Weizenbaum, 1966). People today are often fooled in text-based internet chat rooms and on social networking sites. Sometimes things go wrong, however: Tay, a chatbot developed by Microsoft to interact with 18–24-year-olds on Twitter, learnt to communicate using anonymised public data as well as a group of humans, but rapidly began to tweet like a sexist, racist troll, as well as failing to show any interest in ordinary human topics like pop music or TV.

In 1990, the first annual Loebner Prize competition was held, offering an 18-carat gold medal and a large cash prize for any program that could pass the Turing test, and an annual bronze medal for the most human-like entry of the year. At first no computer came close to passing, despite various restrictions imposed to make the test easier, and Dennett concluded that 'The Turing test is too difficult



ACTIVITY 12.1

A Turing test for creativity

Does it require a conscious human being to paint, draw, or write creatively, or could a machine do as well? If so, it might convince an observer that it was human – in other words it could pass the Turing test. In fact, at least one poem has done just that. As a student, Zachary Scholl used a program which generated poems so convincing that one was accepted for publication without anyone doubting it was written by a human.

For a fun class activity, choose a variety of paintings, pieces of music, jokes, or poems, some of which are created by machines, to see how easily people can tell which is which. A choice of three works well, with people guessing which is machine-created. You can either supply these in advance or ask students to bring in examples without saying where they came from. Get everyone to vote and find out how well the machine creations fare. What kinds of features make people infer human authorship or its absence, and why?

Poetry is a good candidate for this kind of Turing test because poems can be manageably short. See the website for suggestions and for another more classic Turing test activity.



FIGURE 12.7 • The trick, whether you are putting a computer to the Turing test or playing the imitation game, is to know which questions to ask.

AM I A MACHINE?

fool half the judges in a multimodal Turing test including understanding music, speech, pictures, and videos.

Suppose that a machine does pass the test. Suppose that it already has: in 2014 a chatterbot called Eugene Goostman fooled ten out of thirty judges at a Royal Society event in five minutes' conversation during which it pretended to be a 13-year-old Ukrainian boy. What should we conclude about that machine? If the winner were a traditional AI program (as Eugene is), the computational functionalist would conclude that Strong AI had been vindicated, and the program was truly thinking by virtue of running the right program. Other functionalists would argue that such a traditional rule-based program never could pass the test, but that other kinds of machine might, and these would then be truly thinking. Others would insist that whatever the machine is doing, and however well it does it, it is still not *really* thinking like a human does. In other words, it is only pretending to think, or behaving *as if* it is thinking. An alternative is to deny that there is any distinction between 'real' thinking and 'as-if' thinking, a denial that is perhaps in the spirit of Turing's original conception.

The Turing test concerns the ability to think, but all its problems and insights are paralleled in the even trickier question: could a machine be conscious?

COULD A MACHINE BE CONSCIOUS?

Could a machine be conscious? In other words, is there (or could there ever be) 'something it is like to be' a machine? Could there be a world of experience *for* the machine?

'We must be mysterians,' says American philosopher Jesse Prinz. 'The problem isn't that it would be impossible to create a conscious computer. The problem is that we cannot know whether it is possible' (2003, p. 111).

'We have known the answer to this question for a century,' says Searle.

The brain is a machine. *It is a conscious machine.* The brain is a biological machine just as much as the heart and the liver. So of course some machines can think and be conscious. Your brain and mine, for example.

(1997, p. 202)

This sharpens up our question, because what we really mean to ask is whether an *artificial* machine could be conscious; whether we could *make* a conscious

machine. This question is much more difficult than the already difficult question posed by Turing. When he asked 'Can a machine think?', he could cut through arguments about definitions by setting an objective test for thinking.

This just doesn't work for consciousness. First, the arguments about definitions are just as bad, if not worse, because there is no generally agreed definition of consciousness beyond saying that it means subjective experience, or 'what it is like to be' (Chapter 1). Yet many people have a strong intuition that there is nothing arbitrary about it. Either the machine really does feel, really does have experiences, and really does suffer joy and pain, or it does not. This intuition may, of course, be quite wrong, but it stands in the way of dismissing the question 'Can machines be conscious?' as merely a matter of definition.

Second, there is no obvious equivalent of the Turing test for consciousness. If we agree that consciousness is *subjective*, then the only one who can know whether a given machine is conscious is the machine itself, and so there is no sense in looking for an *objective* test.

The problem becomes clearer if you try to invent a test. An enthusiastic robot-builder might, for example, suggest that her machine would count as conscious if it cried when pricked, replied 'yes' when asked whether it was conscious, or pleaded with people not to turn it off. But the sceptic would say, 'It's only got to have an audio recording and a few simple sensors inside it. It's only *pretending* to be conscious. It's a zombie behaving *as if* it's conscious'.

Suppose she decided that it had to laugh at jokes, understand the plot of *Hamlet*, and look deep into your eyes and make you feel loved, and then actually made such a machine. The sceptic might say, 'It's still only *pretending* to be conscious. It's just programmed to respond in human-like ways to jokes and Shakespeare's plays. It doesn't *really* love you.' Then the robot builder might respond, 'But I know it's conscious because if it has a sense of humour, understands human tragedy, and can manipulate human emotions, it *must be*'.

You may have noticed two very familiar arguments here. The robot-builder is a kind of functionalist. She believes that thoughts, beliefs, and subjective experiences are all functional states, so if her robot carries out certain functions then it must be conscious – not because it has some extra mysterious stuff called consciousness that *causes* those things to happen, but because that's what consciousness is. In other words, any machine that could understand *Hamlet*, or look at you in that special way, would necessarily have the linguistic and emotional capacity either to *have* subjective experiences and *be* conscious, or – in illusionist terms – to *claim* to have subjective experiences, and *think* it was conscious.

The sceptic, meanwhile, is a consciousness inessentialist. He believes in zombies. He thinks that however impressive the actions of the machine were, they would not prove it was conscious. His answer would always be, 'It's only pretending. Even if it could do everything you and I do, there would still be nothing it was like to *be* that machine. There would be no light of consciousness on inside'.

If both of these positions are assumed to be legitimate (and they may not be), then there can be no simple test for machine consciousness. Even if functionalists




FIGURE 12.8 • If a robot told you its life story, looked hurt when you offended it, and laughed at your funny stories would you think it was conscious? How could you tell?

'at least one kind of computer can be conscious: the human brain'

(Prinz, 2003, p. 112)

'No one thinks QRio is conscious'

(Greenfield, in Blackmore, 2005, p. 98)

agreed on precisely which functions were the essential ones (which they have not yet done), and designed a test accordingly, the believer in zombies would reject it. Once again (see Chapter 2), believing in zombies seems to lead to an impasse.

Given these difficulties, it might seem impossible to make any progress with the question of machine consciousness, but we should not give up so easily. We may be sure that better and cleverer machines will continue to be built, and that people will keep arguing about whether they are conscious. Even Prinz's (2003) mysterianism is no cause to be defeatist. He urges engineers to keep trying to model minds and learn more about how they work without fooling themselves into thinking they can definitely create conscious machines.

Does it matter? Well, aside from the intellectual quest, there is the problem of suffering – the same problem we faced when thinking about other animals. If machines were conscious, then they could suffer and we, their creators, might need to take some responsibility. This is one of the issues tackled by the field of robot ethics, or roboethics (e.g. Lin, Abney, and Bekey, 2011). Thomas Metzinger asks, 'Should we really try to build conscious machines before we have understood why our own form of subjective experience is accompanied by so much suffering?' (2000, p. 8). Discussing his notion of the phenomenal self model (PSM, Chapter 16), he even suggests that 'we should ban all attempts to create (or even risk the creation of) artificial and postbiotic [partly biological] PSMs from serious academic research' (2003, p. 622). Futurist Ray Kurzweil agrees that the debate over conscious machines lies at the heart of society's legal and moral foundations: 'The debate will change,' he argues,

when a machine – nonbiological intelligence – can persuasively argue on its own that it/he/she has feelings that need to be respected. Once it can do so with a sense of humor [. . .] it is likely that the debate will be won.

(2005, p. 379)

Some dismiss the problem, including Susan Greenfield, who thinks us managing to create machine consciousness is 'so unlikely [. . .] it's like arguing angels on the head of a pin'. If a robot were sent into a burning building to save a person, she would not worry for the robot's sake, 'not for a nanosecond' (in Blackmore, 2005, p. 98). Others are already making plans for dealing with it. For example, in 2007 South Korea began drawing up a Robot Ethics Charter to protect robots from humans and vice versa, and in 2016 the British Standards Institute issued a 'Guide to the ethical design and application of robots and robotic systems'.

In the next section, we will consider some arguments against the possibility of machine consciousness, and in the final section explore ways of making a conscious machine.

CONSCIOUS MACHINES ARE IMPOSSIBLE

There are several plausible – and not so plausible – ways to argue that machines could never be conscious. Some draw on our intuitions about living things and the nature of awareness, and those intuitions can be at once powerful and wrong. It is therefore worth exploring your own intuitions. You may find that some are valuable thinking tools, while others, once exposed, look daft. You may decide

that you want to keep some in spite of the arguments against them, and that with others you want to go through the painful process of rooting them out. Either way, the first step is to recognise them. The story of 'The Seventh Sally' may help (Lem, 1981; see Activity 12.2). Has Trurl just made an amusing model world or has he committed a terrible crime?

Turing (1950) lists nine opinions opposed to his own view that machines *can* think, and some of these are equally applicable to consciousness. Dennett (1995d) and Chalmers (1996) each list four arguments for the impossibility of a conscious robot, and there are many other such lists. Here are some of the main objections to the possibility of conscious machines.

SOULS, SPIRITS, AND SEPARATE MINDS

Consciousness is the unique capacity of the human soul which is given by God to us alone. God would not give a soul to a human-made machine, so machines can never be conscious.

Or you might prefer a nonreligious version of dualism:

Consciousness is the unique capacity the nonphysical mind. We cannot give a separate non-physical mind to a machine, so machines can never be conscious.

Turing strongly disagrees with this argument, and his response is that the builders of thinking machines would not be usurping God's power of creating souls any more than people who have children do: that the builders could be thought of as 'instruments of His will providing mansions for the souls that He creates' (1950, p. 443). The secular equivalent to Turing's riposte would be that if you built the right machine it would automatically attract or create a nonphysical conscious mind to go with it.

If you incline towards the dualist argument in spite of all its difficulties, you might ask yourself the following question. Suppose that one day you meet a truly remarkable machine. It chats to you happily about the weather and your job. It is wonderfully sympathetic when you find yourself pouring out all your emotional troubles. It explains to you, as well as it can, what it feels like to be a machine, and makes



ACTIVITY 12.2

'The Seventh Sally' or How Trurl's perfection led to no good

'The Seventh Sally' is a story from *The Cyberiad* by the Polish writer and philosopher Stanislaw Lem, reprinted with a commentary in Hofstadter and Dennett (1981). We recommend you read the whole story, but here is a brief outline.

Trurl, a brilliant (almost godlike) robotic engineer, or 'constructor', who was well known for his good deeds, wanted to prevent a wicked king from oppressing his poor subjects. So he created an entirely new kingdom for the king. It was full of towns, rivers, mountains, and forests. It had armies, citadels, market places, winter palaces, summer villas, and magnificent steeds, and he 'threw in the necessary handful of traitors, another of heroes, added a pinch of prophets and seers, and one messiah and one great poet each, after which he bent over and set the works in motion'. There were star-gazing astronomers and noisy children, 'And all of this, connected, mounted and ground to precision, fit into a box, and not a very large box, but just the size that could be carried about with ease'. Trurl presented this box to the king, explaining how to work the controls to make proclamations, program wars, or quell rebellions. The king immediately declared a state of emergency, martial law, a curfew, and a special levy.

After a year had passed (which was hardly a minute for Trurl and the king) the king magnanimously abolished one death penalty, lightened the levy, and annulled the state of emergency, 'whereupon a tumultuous cry of gratitude, like the squeaking of tiny mice lifted by their tails, rose up from the box'. Trurl returned home, proud of having made the king happy while saving his real subjects from appalling tyranny.

To his surprise, Trurl's friend was not pleased, but was horrified that Trurl gave the brutal despot a whole civilisation to rule over. But it's only a model, protested Trurl:

all these processes only take place because
I programmed them, and so they aren't

genuine. . . [. . .] these births, loves, acts of heroism, and denunciations are nothing but the minuscule capering of electrons in space, precisely arranged by the skill of my nonlinear craft.

His friend would have none of it. The size of the tiny people is immaterial, he said, 'And don't they suffer, don't they know the burden of labor, don't they die? [. . .] And if I were to look inside your head, I would also see nothing but electrons'. Trurl, he says, has committed a terrible crime. He has not just imitated suffering, as he intended, but has created it.

What do you think?

For a group discussion

This story can provoke heated and insightful disagreements. Ask everyone to read the story in advance and to write down their answer to the question 'Has Trurl committed a terrible crime?': 'Yes' or 'No'. Check that they have done so, or ask for a vote.

Ask for two volunteers with strong opinions, one to defend Trurl, the other to accuse him of cruelty. This works best if the participants really believe in their respective roles. Others can ask questions and then vote. Has anyone changed their mind? If so, why? Is there any way of finding out who is right?

you laugh with funny stories about humans. Now what do you conclude?

- 1 the machine is a zombie (with all the familiar problems that entails)
- 2 God saw fit to give this wonderful machine a soul or, if you prefer, the machine had attracted or created a separate mind
- 3 you were wrong, and a machine can be conscious.

This is a good thought experiment for winking out implicit assumptions and strongly held intuitions. Turing suggests that fear and a desire for human superiority motivate the theological objection and also what he calls the 'Heads in the Sand' objection: 'The consequences of machines thinking would be too dreadful. Let us hope and believe that they cannot do so' (1950, p. 444). Some people may similarly fear the possibility of a machine being conscious.

THE IMPORTANCE OF BIOLOGY

Only living, biological creatures can be conscious, therefore a machine, which is manufactured and non-biological, cannot be.

At its simplest, this argument is mere dogmatic assertion, or an appeal to vitalism. Yet it might be valid if there were shown to be relevant differences between living and non-living things. For example, it might turn out that only protein membranes just like those in real neurons can integrate enough information, quickly enough and in a small enough space, to make a conscious machine possible, or that only the neurotransmitters dopamine and serotonin can sustain the subtlety of emotional response needed for consciousness. But in this case,

robot-builders would probably make use of these chemicals, overcoming the objection by blurring the distinction between natural and artificial machines. There are already robots that feed on flies and slugs, and people who have heart valves, cochlear implants, prosthetic limbs, and 'neuroprostheses', so this is far from science fiction.

A second argument is that biological creatures grow and learn over a long period before they become conscious; machines have no history and so cannot be conscious. This has some force if you think only of machines made in factories and pumped out ready to go, but perhaps the best (or only) way to make effective robots is to give them time to learn in a real environment. It is clear from connectionism, embodied cognition, and situated and swarm robotics that such periods of environmentally embedded learning may well be necessary.

Searle (1992) claims that 'brains cause minds', and that there is something special about biology. His theory of 'biological naturalism' seems to imply that brains and minds must be distinct from each




FIGURE 12.9 • Cog, MIT's upper-body anthropomorphic robot, interacting with its technician.

other, but he denies being a property dualist or indeed any kind of dualist (2002). He stresses that although consciousness is *causally* reducible to its neurological base, it is *ontologically* distinct from the brain in the sense that it must be experienced (Chapter 17). He explains that ‘biological brains have a remarkable biological capacity to produce experiences, and these experiences only exist when they are felt by some human or animal agent’ (1997, p. 212). Even so, Searle does not claim that brain tissue is necessary for consciousness. He argues that other systems could be conscious too, but only if they had equivalent causal powers to those of the brain. However, he does not say what those causal powers are.

MACHINES WILL NEVER DO X

There are some things that no machine can possibly do because those things require the power of consciousness.

Turing (1950, p. 447) offers a selection of things said to be impossible for a machine:

Be kind, resourceful, beautiful, friendly [. . .] have initiative, have a sense of humor, tell right from wrong, make mistakes [. . .] fall in love, enjoy strawberries and cream [. . .] make someone fall in love with it, learn from experience [. . .] use words properly, be the subject of its own thought [. . .] do something really new.

It is a good list; nearly seventy years later machines still cannot do most of them. Yet, as Turing points out, the claim is based on people’s extrapolation from machines they have actually seen, rather than any principled reason why machines could not do such things. People too easily jump to several conclusions: first, that machines cannot do X; second, that because we can do X, we must have something machines cannot have; and third, that this extra thing is consciousness.

The last is particularly interesting and relates to what is often called ‘Lady Lovelace’s objection’. Ada Lovelace, Lord Byron’s daughter, studied mathematics, became fascinated by Babbage’s ideas, and wrote the only full account we have of his Analytical Engine. She famously said that ‘The Analytical Engine has no pretensions to *originate* anything. It can do *whatever we know how to order*



CONCEPT 12.2

HUMANOID ROBOTS AND SIMULATIONS

Robots that look and move like people have been built for fun, companionship, and domestic tasks as well as for research. Sony’s Qrio and WowWee’s RoboSapien were meant for entertainment, while Honda’s Asimo and the ‘social robot’ Nadine were intended to help people and ultimately to act as carers. Able to walk and carry things, they have limited cognitive abilities, including recognising people and remembering and responding to speech, but make no claims to subjective experience. Projects that copy aspects of human functioning to investigate consciousness take three main approaches.

1 Humanoid robots

The ambitious Cog project, conceived in 1993 by Brooks and colleagues at MIT (Brooks et al., 1998), aimed to learn about human cognition by trying to implement it. Consisting of a humanoid body with arms and a head, it had dozens of motors, a core of hundreds of connected PCs, moveable eyes, and integrated auditory, vestibular, and tactile sensory systems.

Built on the principles of embodied cognition, Cog was given no detailed internal representations of the world but learned through the coupling between its own actions and perceptions. Its basic social skills included finding faces and eyes, sharing attention with a human, following pointing gestures, and imitating head nods. An original aim, that Cog would acquire the cognitive abilities of a young child, was never achieved, but there were many surprises, including the way its carers started treating Cog as if it *mattered* what they did to it. This was even more obvious with the ‘sociable robot’ Kismet, which was just a moveable head with large eyes and moving red lips, with simple routines designed to be cheap, fast, and just adequate. It could move towards or away from things to suit its cameras, giving the impression of being interested in things and people, and had three ‘mood’ variables expressed by making sounds and changing its

facial expression. The interesting thing was how people responded to Kismet, talking, coaxing, and mirroring its facial expressions. They behaved as though Kismet was alive and had feelings.

But could things *really* matter to Cog or Kismet? This depends what you think about *real* mattering and *real* suffering: are they special biological or human attributes forever denied to machines, or are they just more of the kind of thing these primitive robots already have (Dennett 1998b)?

2 Anthropomimetic robots

Rather than imitating the appearance of a human, Owen Holland's CRONOS robot imitates aspects of the functioning of a human body in 'a strongly embodied approach to machine consciousness' (Holland, 2007). Holland wanted to embody in a robot the kinds of inner representations that would allow it to interact with the real world in human-like ways. CRONOS is an upper-body robot with plastic bone-like structures, elastic tendons, and functional joints, making it wobbly and hard to control, but it allowed its creators to test various internal models of world and self that, they argue, provide insight into human consciousness and self-consciousness. CRONOS was subsequently developed into the more ambitious ECCE robot, which continues the search for human-like mentality in an anthropomimetic machine.

3 Disembodied simulation

Existing entirely as programs inside computer systems, Rodney Cotterill's CyberChild is a simulation of a baby based on the known circuitry of mammalian nervous systems. With just two senses, hearing and touch, CyberChild can feel pain when its simulated stomach is empty or its nappy unchanged, and can feed, cry, and control its simulated bladder. In other words, it's a crying baby simulated in silicon. Cotterill hopes to be able to infer the presence of consciousness from its behaviour, and even to discover the NCCs this way. He claims 'there appears to be no fundamental reason why consciousness could not ultimately develop and be observed' (2003, p. 31). Not everyone would agree.

it to perform' (in Taylor, 1843, p. 772f) This suggests that the machine could not be creative, and the same argument has often been applied to modern computers: 'Working in a fully automated mode, [computers] cannot exhibit creativity, emotions, or free will. A computer, like a washing machine, is a slave operated by its components' (Buttazzo, 2001, p. 26). But this argument seems less and less applicable as time goes by. Computers can already write poems, essays, and screenplays, make pictures and compose music. Some do this by simple algorithms combining ready-made segments, some use neural networks and parallel processing, and some use genetic or evolutionary algorithms.

These are computer versions of the familiar evolutionary algorithm (Chapter 10): they 1) take a segment of computer code or program, 2) copy it with variations, 3) select from the variants according to some specified outcome, and 4) take the selected variant and repeat the process. Is this *real* creativity or only *as-if* creativity? That rather depends on what you think *real* creativity is. If you think that real creativity requires some special power of consciousness, then perhaps machines are not really creative. But what if human creativity depends on the evolutionary algorithm, using exactly the same processes as those described above? This would mean the copying, selection, and recombination of old memes to make new ones. In that case, biological creativity, human creativity, and machine creativity would all be examples of the same evolutionary process in operation and none would be more *real* than the others (Blackmore, 2007b).

A variant on this argument is what Turing calls the 'mathematical objection'. There are some things that machines cannot do, so if we can do any of them that proves we have something extra – consciousness.

As we have seen (Concept 12.1), there are some functions that are noncomputable, meaning there are questions a machine could never answer correctly, however much time it is given (Turing, 1950). Penrose (1989, 1994a) claims that mathematicians can intuitively see noncomputable truths, and this *real* understanding requires conscious awareness. So consciousness itself must be beyond computation.

This is why he thinks we need an entirely new kind of physics to understand consciousness, and so proposes the theory of objective reduction in the microtubules (Chapter 5).

Kurzweil (1999) retaliates that 'It is true that machines can't solve Gödelian impossible problems. But humans can't solve them either' (p. 117). We can only estimate them, and so can computers, including quantum computers. And Turing himself pointed out that we humans are notoriously error-prone and might even revel in our limitations. Could machines revel in their limitations? asks Hofstadter (2007): could a machine be confused? Could it know it was confused?

We do not know the answer, but it seems that none of these arguments proves the impossibility of building a conscious machine.

If there are some things that machines can never do, we are far from knowing what they are and why.



PRACTICE 12.2

IS THIS MACHINE CONSCIOUS?

As many times as you can, every day, ask: '***Is this machine conscious?***'

This exercise, like the one about animal consciousness, is directed out beyond yourself. Whenever you use a phone, laptop, or TV, or depend on air traffic control or satellite navigations systems, ask 'Is this machine conscious?' You can do the same with fridges, cars, and electronic games, or indeed anything you like. Explore your own intuitions. Can you discern the reasons why you are more tempted to attribute some inkling of consciousness to some machines rather than others?

None of the general arguments considered so far has demonstrated that a machine cannot be conscious. Two further arguments are much more specific and much more contentious.

THE CHINESE ROOM


Among Turing's list of arguments against machine thinking is 'The argument from consciousness'. This, he says, might be used to invalidate his test, because 'the only way by which one could be sure that a machine thinks is to *be* the machine and to feel oneself thinking' (Turing, 1950, p. 446). Even if the machine described its feelings, we should take no notice. He rejects this argument on the grounds that it leads only to solipsism – the view that we can never know anything about other minds than our own – and in this way defends his test. Yet this argument was not to be so easily defeated. Thirty years later, it gained its most powerful advocate in the philosopher John Searle, with his famous Chinese Room thought experiment.

Searle proposed the Chinese Room as a refutation of Strong AI – that is, the claim that implementing the right program is all that is needed for understanding. It is most often used to discuss intentionality and meaning with respect to AI, but many people, including Searle himself, believe that the Chinese Room has



PROFILE 12.2

John Searle (b. 1932)



John Searle is Professor Emeritus of Philosophy at Berkeley, where he has been since 1959. He says he is, and always has been, ‘interested in everything’. As a student at the University of Wisconsin, he resigned as President of the student board to study harder. He then went to Oxford, where he spent three years as a Rhodes scholar and became a don at Christ Church. He has written books on language, rationality, and consciousness, including *The Rediscovery of the Mind* and *Seeing Things as They Are: A Theory of Perception*, and his books have been translated into more than twenty languages. His Chinese Room thought experiment is probably the best-known argument against the possibility of ‘Strong AI’, a term that he invented. He says that ‘brains cause minds’ and argues for ‘biological naturalism’.

important implications for consciousness. In an echo of one of the criteria for animal consciousness we considered in Chapter 10, it makes language central.

Searle took as his starting point Roger Schank’s programs that used scripts to answer questions about ordinary human situations, such as having a meal in a restaurant. These were firmly in the GOFAL tradition, manipulating symbols according to formal rules, and incorporating representations of relevant knowledge. Supporters of strong AI claimed that these programs really understood the questions and their answers. This is what Searle attacked.

‘Suppose that I’m locked in a room and given a large batch of Chinese writing. Suppose furthermore (as is indeed the case) that I know no Chinese, either written or spoken’, begins Searle (1980, p. 417–418). Inside his room, Searle has lots of Chinese ‘squiggles’ and ‘squoggles’, together with a rule book in English. People outside the room pass in two batches of Chinese writing which are, unbeknown to Searle, a story, in Chinese of course, and some questions about the story. The rule book tells Searle which squiggles and which squoggles to send back in response to which ‘questions’. After a while he gets so good at following the instructions that from the point of view of someone outside the room his ‘answers’ are as good as those of a native Chinese speaker. He next supposes that the outsiders give him a story and questions in English, which he answers these just as a native English speaker would – because he is a native English speaker. So his answers in both cases are indistinguishable. But there is a crucial difference.

In the case of the English stories, he *really* understands them. In the case of the Chinese stories, he understands nothing.

So here we have John Searle, locked in his room, acting just like a computer running its program. He has inputs and outputs, and the rule book to manipulate the symbols, but he does not understand the Chinese stories. The moral of the tale is this: a computer running a program about Chinese stories understands nothing of those stories, whether in English or Chinese or any other language, because Searle has everything a computer has, and he does not understand Chinese.

Searle concludes that whatever purely formal principles you put into a computer, they will not be sufficient for *real* understanding. Another way of putting it is that you cannot get semantics (meaning) from syntax (rules for symbol manipulation). Any meaning or reference that the computer program has is in the eye of the user, not in the computer or its program. So Strong AI is false.

The Turing test is also challenged because in both languages Searle claims he passes the test perfectly, but in English he *really* understands while in Chinese he doesn’t. Note that, for Searle, this shows that there is something extra that he has and the computer does not. This something is *real* (as opposed to *as-if*) intentionality (the capacity to *be about* something). He concludes that ‘Whatever it is that

‘no program [...] is sufficient for intentionality’

(Searle, 1980, p. 424)

the brain does to produce intentionality, it cannot consist in instantiating a program since no program, by itself, is sufficient for intentionality' (1980, p. 424). The something is also, he claims, *subjective*, and this is where the argument becomes directly relevant to consciousness.

Reaction to the Chinese Room has been ferocious for decades. Searle (1980) himself listed six replies and rebutted them in turn, and many more followed. Among them, the 'systems reply' argues that while Searle himself might not understand Chinese, the whole distributed cognitive system consisting of him and the room does. Searle responds that he could internalise all the rules and do the manipulations in his head and he still wouldn't understand Chinese. The 'robot reply' suggests putting a computer into a robot and letting that interact with the outsiders, claiming that a machine which could interact with the world the language refers to *would* understand, but Searle responds that adding a set of causal relations with the outside world makes no difference because you could put him inside the robot and he would still just be manipulating symbols and would still not understand Chinese. The 'brain simulator reply' proposes a program that simulates the actual sequence of neuron firings in a real Chinese brain. Searle responds that as long as this program only simulates the formal properties of the brain, it misses the crucial causal properties that allow brains to cause minds: the properties that cause consciousness and intentional states.

The argument started as a refutation of Strong AI. Have things changed with the advent of connectionism and behaviour-based robotics? The robot reply was a step in this direction because it suggested that interaction with the real world was essential for understanding or intentionality. As McGinn puts it, 'Internal manipulations don't determine reference, but causal relations to the environment might' (1987, p. 286). Another way of saying this is that the symbols must be grounded in the real world, because it is only through symbol grounding that we humans come to understand and have intentional states (Harnad, 1990; Velmans, 2000). Similarly, Chalmers (1996) points out that a computer program is a purely abstract object, while human beings are physically embodied and interact causally with other physical objects. The bridge between the abstract and the concrete, he says, lies in *implementation*. Having the right program is not sufficient for consciousness, but implementing it is. Ron Chrisley (2009) promotes a 'moderate AI' position: that modelling necessarily uses properties shared by AI systems and brains, but instantiating these common properties is not sufficient for consciousness. Something more, such as symbol grounding or biology, might be needed. Moderate AI, he says, is immune to the Chinese Room argument.

Dennett presses a version of the systems reply. The problem with this thought experiment, he suggests, is that Searle misdirects our imagination by luring us




FIGURE 12.10 • Searle asks us to imagine that he is locked in a room. People pass in squiggles and squaggles. He looks up what to do in his rule book and passes out more squiggles and squaggles. Unbeknown to him the symbols being passed in are Chinese stories and questions, and the symbols he passes out are answers. To the people outside he seems to understand Chinese, but he is like a computer, manipulating symbols according to rules, and he does not understand a word.

Searle's Chinese Room article is a 'religious diatribe against AI, masquerading as a serious scientific argument'

(Hofstadter, in Searle, 1980, p. 433)

into imagining that a very simple table-lookup program could do the job, when really 'no such program could produce the sorts of results that would pass the Turing test, as advertised' (Dennett, 1991, p. 439). Complexity does matter – so even if a hand calculator does not understand what it is doing, a more complex system, like one that passes the Turing test, could. He suggests that we should think of understanding as a property that emerges from lots of distributed quasi-understandings in a large system (p. 439).

We might go even further and reject Searle's thought experiment (like the zombie argument or Mary the colour scientist we considered in Chapter 2) on the grounds that it instructs us to imagine something impossible. Searle claims that with only the Chinese symbols and his rule book (or even with the rules memorised and inside his head), he really could pass the Turing test without understanding a word of Chinese. But what if he couldn't? It might turn out that symbol grounding, or learning by interactions with the real world, or something else again, is necessary for passing the test as well as for 'really understanding' a language. In this case, there are only two options. Either he does not have these necessities, and his symbol manipulations fail to convince the Chinese people outside, or he does, and that means he comes to understand Chinese in the process. Either way, the scenario Searle described in the original thought experiment might be impossible.

Just as with Mary and zombies, there is no final consensus on what, if anything, the Chinese Room shows. Some people think it shows nothing. Some people think it demonstrates that you cannot get semantics from syntax alone, and that a machine could not be conscious simply by virtue of running the right program. Some (perhaps a minority) agree with Searle that it demonstrates a fundamental difference between the *real, conscious* intentionality that we humans have, and merely *as-if* intentionality. In this case, machines could be conscious only if they had the same causal properties as living human brains, whatever those properties are.

IS THIS MACHINE CONSCIOUS?

HOW TO BUILD A CONSCIOUS MACHINE

Many roboticists and computer engineers ignore all the arguments and simply get on with pursuing their 'Holy Grail': 'the artificial consciousness quest – nothing less than the design of an artificial subject' (Chella and Manzotti, 2007, p. 10). There are two main ways of setting about the task. The first asks how to build a machine that *seems* to be conscious; the second asks how to build a machine that really *is* conscious (whatever that means).

But some say there is no need for a grand quest, for conscious artificial machines are all around us already.

THEY'RE ALREADY CONSCIOUS

In 1979, John McCarthy, one of the founders of AI, claimed that machines as simple as thermostats can be said to have beliefs. John Searle was quick to challenge him, asking 'John, what beliefs does your thermostat have?' Searle

'My thermostat has three beliefs – it's too hot in here, it's too cold in here, and it's just right in here.'

(McCarthy, in Searle, 1984, p. 30)

admired McCarthy's courageous answer, for he replied, 'My thermostat has three beliefs – it's too hot in here, it's too cold in here, and it's just right in here' (Searle, 1984, p. 30).

The thermostat was an unfortunate choice for Searle but lucky for McCarthy. Although extremely simple, it has two of the crucial features required of an autonomous agent: it perceives its environment, and it responds to changes by acting on that environment. A thermostat is not an abstraction or a disembodied computation; it is grounded in the real world through its actions – simple as they are.

You might think that McCarthy was joking, or that he didn't mean that thermostats have *real* beliefs like ours. But this implies that you think there is a difference between *real* intentionality and only *as-if* intentionality. Do you? As we've seen, Searle argues that only biological human beings have the *real* thing, whereas computers and robots behave *as if* they understand languages, believe things, and have experiences. If you agree with Searle, then you have to decide what the difference is between the real thing and the simulation.

If you reject the distinction, you might say that the beliefs of thermostats are just as real as human beliefs, although far simpler, or you might say that the whole idea of *real* beliefs is misguided, and that all human intentionality is *as-if* intentionality. Either way, humans and machines have the same kind of beliefs and we humans are already surrounded by believing machines.

Are any of today's machines conscious? To some people, intentionality (being about something) entails or requires consciousness. The Chinese Room argument was designed to deal with intentionality, but both Searle and some of his critics applied it to consciousness, in the sense that only a conscious being could *really* understand Chinese. On this interpretation, if any machine has beliefs (one kind of intentionality), it must thereby be conscious.

Others distinguish consciousness from intentionality, but then the same dichotomy between *real* and *as-if* arises for consciousness too. If you think this way, then robot-builders need to find out what *real* consciousness is and whether a machine could have it. Alternatively, if there is no difference between *real* consciousness and *as-if* consciousness, we humans are already sharing our world with the beginnings of AC.

I have no more hope, nor project, nor strength, nor will, I go and I live like a wheel that has been pushed and that will roll until it falls over, like a leaf that flies on the wind as long as the air holds it up, like the thrown stone that falls until it finds the bottom – a human machine that sheds tears and secretes pain, an inert thing that finds itself here without cause, created by an incomprehensible force and understanding nothing about itself.

(Gustave Flaubert, *Sentimental Education*
[L'Éducation sentimentale], 1869)

'Room thermostats are not conscious'

(Aleksander, 2007, p. 97)

'What is your extra ingredient, and why should that account for conscious experience?'

(Chalmers, 1995a, p. 207)

FIND X AND PUT IT IN A MACHINE

Suppose that humans have some magic ingredient 'X', by virtue of which they are *really* conscious. If we wanted to make a conscious machine, we might then proceed by finding out what X is and putting it into a machine, or we might build a machine in such a way that X would naturally emerge. The machine would then, theoretically at least, be conscious.

Chalmers (1995a) says that those who are serious about solving the hard problem need to find the right 'extra ingredient' to account for conscious experience. McGinn (1999) calls the property that would explain consciousness C*, and asks whether C* is possible in inorganic materials or not. He concludes that we cannot possibly know. According to his mysterian theory the human intellect is incapable of understanding how organic brains become conscious, so there is no hope of us ever finding C* or knowing whether a machine could have it.

Others are less pessimistic. British AI researchers Aaron Sloman and Ron Chrisley are not deterred in their search for machine consciousness by the fact that *'We do not yet have the concepts necessary for fully understanding what the problem of consciousness is'* (2003, p. 140). One of the strongest proponents of AC is David Chalmers, who rejects the Chinese Room and other arguments against computationalism. Even though he is a dualist of sorts, he claims that any system with the right sort of functional organisation would be conscious. He argues 'not just that implementing the right computation suffices for consciousness, but that implementing the right computation suffices for rich conscious experience like our own' (1996, p. 315). He does not go on to say what 'the right computation' is, but he has defended a very broad notion of computation as a foundation for artificial intelligence, claiming that in it 'the causal structure of mentality is replicated' (1993/2011). So, Chalmers suggests trying to find X as a way forward.

'a model that is computationally equivalent to a mind will itself be a mind'

(Chalmers, 1993/2011)

How might we do this? One way is to make a list of criteria for a conscious machine: a list of possible Xs. Philosopher Susan Stuart (2007) suggests 'engaged embodiment, goal-directed animation, perception and imagination' and the ability to synthesise experiences and recognise them as its own experiences, and she emphasises the importance of kinaesthetic as well as cognitive imagination.

AI researcher Igor Aleksander tackles phenomenology 'as the sense of self in a perceptual world' and starts from his own introspection to break this down into five key components or axioms (Aleksander and Morton, 2007). He then uses these as criteria for a conscious machine (Aleksander, 2007). They are:

- 1 Perception of oneself in an 'out there' world.
- 2 Imagination of past events and fiction.
- 3 Inner and outer attention.
- 4 Volition and planning.
- 5 Emotion.




FIGURE 12.11 • Can we find X and put it in a machine?

On this basis, Aleksander develops an abstract architecture called the Kernel Architecture (KA) that incorporates all five. A key mechanism is depiction: a direct

representation of where elements of the world are that allows attention to be directed appropriately. As Aleksander and Morton point out, there are known to be many cells doing this kind of job in the human brain. Importantly, KA also includes depiction of the self in its world. Aleksander concludes that a robot might be said to be conscious if equipped with KA or in some way manages to model both itself and the world. He adds that on his model higher-order thought theories turn out to be about our ability to attend to different parts of the architecture that allows this and translate their activity into language.

Among others who begin from basic principles is British physicist John Taylor, whose attention-based CODAM model takes principles of human attention as a starting point for building a conscious machine with a sense of self.

A different approach is to start from existing theories of consciousness and build machines that implement them. For example, according to global workspace theories, the contents of consciousness are whatever is being processed in the GW. The GW is itself a large network of interconnected neurons, and its contents are conscious by virtue of the fact that they are made globally available to the rest of the system, which is unconscious. On these theories, 'X' is global availability. So, presumably, a machine should be conscious if it is designed with a GW whose contents are made available to the rest of its system.

American mathematician Stan Franklin (2003) built a software agent called IDA. This 'Intelligent Distribution Agent' was developed for the US Navy to help solve the problem of assigning thousands of sailors to different jobs. To do this, she has to communicate with the sailors by email in natural language, as well as satisfying numerous Navy policies and job requirements. IDA is built on the basis of GW architecture, with coalitions of unconscious processes finding their way into a global workspace from where messages are broadcast to recruit other processors to help solve the current problem. Franklin describes IDA as being *functionally* conscious in the sense that she implements much of GWT, but not phenomenally conscious or self-conscious, although he argues that building in a simple kind of self based on Damasio's ideas of the proto-self would be quite feasible.

IDA has since been developed into LIDA (Learning IDA), who is capable of perceptual, episodic, and procedural learning (Franklin and Patterson, 2006). Baars and Franklin argue that the functions of consciousness are produced by adaptive, biological algorithms and that 'machine consciousness may be produced by similar adaptive algorithms running on the machine' (2009, p. 23). Since LIDA implements much of the functionality of GWT, they conclude that she may be 'functionally conscious'. They also suggest that she could one day be made phenomenally consciousness by adding mechanisms that produce perceptual stability or that implement various notions of self in a LIDA-controlled robot (Baars and Franklin, 2009).

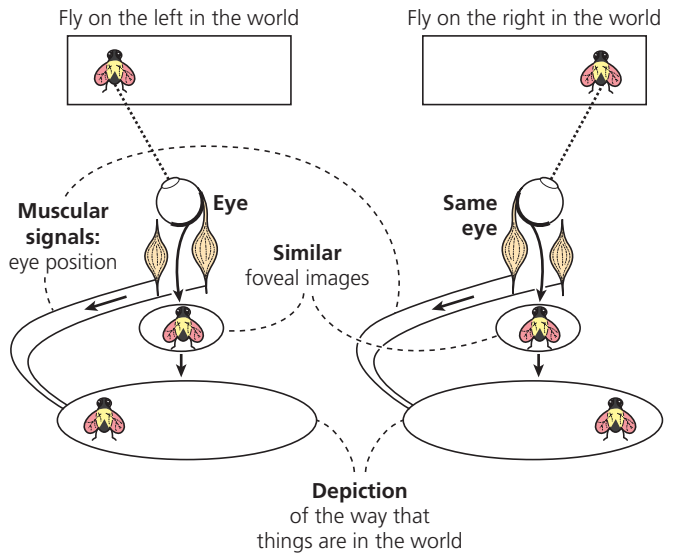


FIGURE 12.12 • An example of a depictive process which involves muscle action. A depiction arises in an area where the foveal image is 'positioned' by the muscular signals that indicate eye position (after Aleksander, 2005, p. 39).



PROFILE 12.3

Owen Holland (b. 1947)



Owen Holland is best known for his work on machine consciousness and for building biologically inspired robots, but he only started robotics as a hobby in 1988 after working as a production engineer, boatbuilder, transport manager,

insurance salesman, and chef in a steak bar. He had a croft in Orkney for eight years where he built his own house and tended cows, goats, ducks, and chickens, grew oats, and made hay. Just as eclectically, he has had academic positions in psychology, electrical engineering, computer science, and cognitive robotics at universities in England, Scotland, Germany, Switzerland, and the USA. He worked on two robot projects at Caltech and helped set up the robotics lab at the University of the West of England, Bristol. Holland used the biologically inspired robot CRONOS to ask whether it could be phenomenally conscious according to various theories of consciousness. CRONOS has since developed into the anthropomimetic ECCE robot which, with its human-like structures, may lead to human-like cognition. He is now Emeritus Professor of Cognitive Robotics in the Sackler Centre for Consciousness Science at the University of Sussex.

‘Engineering will step from the mere design of complex artefacts to the design of subjects’

(Chella and Manzotti, 2007, p. 11)

Note that IDA and LIDA are software agents and so, like KA, are not permanently tied to any particular physical machine, raising the question of just what it is that we think might be conscious. Could a complicated mass of software or a virtual machine without a single material body be conscious? Could it be that self-representations rather than people (or bats, or octopuses) are the subjects of experience? If so, there might already be conscious entities living in cyberspace and supported by multiple machines in different locations. Come to think of it, when we refer to ourselves as conscious, are we referring to our bodies, our brains, our inner selves, or something else entirely? This is another interesting conundrum that MMC may throw light on.

A strongly embodied example is Owen Holland’s CRONOS, an anthropomimetic upper-body robot designed to include internal models of itself and the world (Holland, 2007; Holland, Knight, and Newcombe, 2007). The idea of internal models began over sixty years ago when Craik suggested that intelligent organisms might need small-scale models of external reality and of their own possible actions. For a long time Holland rejected this idea and developed purely behaviour-based robotics with no internal models, but he later returned to this principle in creating CRONOS and its successor ECCE Robot. With human-like elastic muscles and tendons, and a bone-like skeleton, CRONOS has a single colour camera for an eye, an elongated neck to help it inspect objects, and complex moving arms. It builds models of the world around it by moving its eye and looking at and interacting with objects, and uses a model of its own body and capabilities to plan its possible actions. It does not interact with people, however, and has no language or emotions.

How does this relate to consciousness? The robot has a model of itself, the Internal Agent Model (IAM), and a model of the world, the Internal World Model (IWM), and it uses these to act and to track changes in its own body and beyond. These

internal simulations form the basis for what its creators call ‘functional imagination’: the ability to manipulate information that is not directly available to sensors (Marques and Holland, 2009). These models may not be completely detailed and accurate, but they are the only self and world that the robot knows. ‘Unless it somehow had access to the information that it was only a model, the IAM would operate as if it were the agent itself’ (Holland, 2007, p. 101). Critically, this means that the IAM would be transparent in the sense used by Metzinger (2003, 2009) when describing his phenomenal self-model: the robot depends on a model which doesn’t include the fact that it’s a model. This model would in some sense describe itself as an embodied agent – rather as we do.

Gauges in the head, Szpindel had called them. But there were other things in there too. There was a model of the world, and we didn’t

look outward at all; our conscious selves saw only the simulation in our heads, an interpretation of reality, endlessly refreshed by input from the senses. What happens when those senses go dark, but the model – thrown off-kilter by some trauma or tumor – fails to refresh? How long do we stare in at that obsolete rendering, recycling and massaging the same old data in a desperate, subconscious act of utterly honest denial? How long before it dawns on us that the world we see no longer reflects the world we inhabit, that we are blind?

(Peter Watts, *Blindsight*, 2006, p. 193)

Some other examples of theory-led robotics might include Giulio Tononi's (2015) integrated information theory (if you build integration into a machine in ways that increase Φ , the corresponding degree of consciousness should follow) or Michael Graziano's attention schema theory (if a system can model its own attention, it can lay claim to consciousness) (Webb and Graziano, 2015).

According to quantum theories, none of these implementations would produce *real* consciousness because that needs quantum processes. For example, in Penrose and Hameroff's version, consciousness emerges from quantum coherence in the microtubules, so one would need to build a quantum computer that achieved this kind of integration across its system. One might then conclude that it was *really* conscious.

None of this avoids the two big problems mentioned at the start of this section. First, we do not know what consciousness is. Each of these theories (and many others) says something about what consciousness is or what it emerges from, but if the appropriate machine were built, critics could still argue that this particular theory was wrong and therefore the machine was not conscious after all. Second, we have no test for proving whether a machine is conscious or not, so even if one of these machines claimed to be conscious, stayed awake all night worrying about consciousness, and passed the Turing test, we could still not convince sceptics that it was *really* conscious, even though we might have learned a lot from the machine.

DELUDED MACHINES

There is a completely different way of thinking about X. Perhaps consciousness is not what it seems to be, and we are in some fundamental way deluded about the nature of consciousness. According to this view, we may believe we are conscious observers, experiencing a continuous stream of contents passing through our conscious minds, but we are wrong because there is no Cartesian Theatre, no audience, no 'actual phenomenology', and no continuous stream of conscious experiences (Dennett, 1991; Blackmore, 2002, 2012). We humans certainly *seem* to be conscious, and that requires explaining, but the right kind of explanation




FIGURE 12.13 • CRONOS is a strongly embodied anthropomorphic upper-body robot, with human-like elastic muscles and tendons, and bone-like skeleton.




FIGURE 12.14 • The large square represents the agent, and the large circle the world. The thought bubbles represent the agent's internal models of itself and the world, which are separate but which interact to give functionally useful predictions of the effects of possible actions (after Holland, 2007).

is one that accounts for why we have this particular illusion. This means that a machine would have human-like consciousness only if it were subject to the same kind of illusion. The task is then to understand how the illusion comes about and design a similarly deluded machine.

One possible example is Sloman and Chrisley's (2003) CogAff architecture, developed as a framework for thinking about both natural and artificial information processing systems, and based on the implicit theory that minds are information-processing virtual machines. They propose 'virtual machine functionalism' (VMF), which avoids some of the problems of other forms of functionalism by including internal processes of the virtual machine that do not have to be closely linked to its input-output relations.

The CogAff architecture can be structured in various ways, for example having a 'multi-window' perception and action system as opposed to restricting the routes through it to give 'peep-hole' perception and action. Or it can use subsumption architecture that includes a deliberative reasoning (what-if) layer as well as the reactive layer. Critically, it also has a meta-management layer that allows it to attend to aspects of its own internal functioning.

But what about the 'qualia – the private, ineffable way things seem to us'? Sloman and Chrisley want to 'explain qualia by providing an explanation of *the phenomena that generate philosophical thinking of the sort found in discussions of qualia*' (2003, p. 165). Their virtual machine includes processes that classify its own internal states. Unlike words that describe common experiences (such as seeing 'red' in the world), these refer to internal states or concepts that are not strictly comparable from one virtual machine to another – just like qualia. If people protest that there is 'something missing' – the indefinable quality, the what it's like to be, or what zombies lack – their reply is that the fact that people think this way is part of what needs to be explained, and their approach can do just that.

Extending this line of thinking, another obvious contributor to illusions is language. For example, the self has been described as a construct of language, a 'center of narrative gravity', a 'benign user illusion' that emerges in creatures who use language, or a 'selfplex' constructed by and for the replication of memes (Chapters 5 and 11). The implication here is that if any machine – or indeed any non-human animal – were capable of using language, and capable of talking about 'I', 'me', and 'mine', it would also fall for the illusion that it was an experiencing self, and would then be conscious like us.

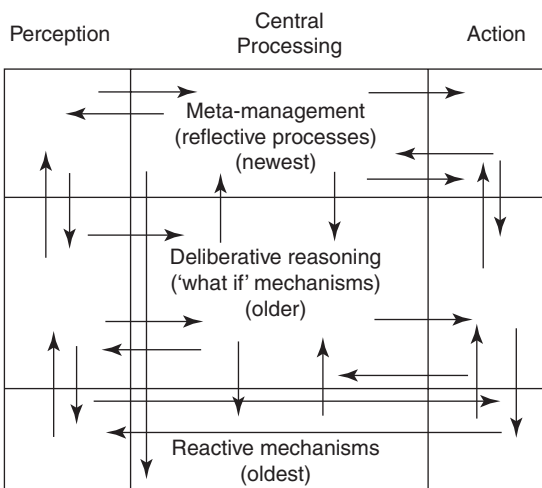


FIGURE 12.15 • The CogAff schema: superimposing towers and layers (Sloman and Chrisley, 2003, p. 163).

SPEAKING MACHINES

Charles Darwin's grandfather, Erasmus Darwin, built a speaking machine around 1770 that could (just about) say 'Mama' and 'Papa'. An entrepreneur offered him £1,000 if he could get it to recite the Lord's Prayer and the Ten Commandments, but his money was safe. Since then we have enjoyed fantastic

speaking machines that play recorded speech, read aloud from printed text, or turn spoken language into print. Then there are computers that will tell you, in a perfectly comprehensible, if annoying, voice, that they think you have made a

mistake and would like you to turn around when possible. None of these, however, can probably be said to understand what they say.

Early attempts to teach machines language used the GOFAL approach, trying to program computers with the right rules. But natural languages are notoriously resistant to being captured by rules of any kind. Such rules as there are always have exceptions, words have numerous different meanings, and many sentences are ambiguous. A machine programmed to parse a sentence, construct a tree of possible meanings, and choose the most likely may completely fail on sentences that you and I have no trouble understanding. Pinker (1994, p. 209) gives some examples:

Ingres enjoyed painting his models in the nude.

My son has grown another foot.

Visiting relatives can be boring.

I saw the man with the binoculars.

The most famous example was encountered by an early computer-parser in the 1960s. The computer came up with no less than five possible meanings for the well-known saying 'Time flies like an arrow', giving rise to the aphorism 'Time flies like an arrow; fruit flies like a banana'.

Machines analysing language this way remained like Searle inside his Chinese Room, shuffling symbols back and forth. The advent of neural nets and connectionism improved the prospects. For example, early neural nets learned relatively easily how to pronounce written sentences correctly without being programmed to do so, even though the correct pronunciation of a word often depends on the context. Even so, they could not be said to speak or understand true language.

A real shift occurred with an approach that is closer to evolutionary theory and memetics. One of the fundamental principles in memetics is that when organisms can imitate each other, a new evolutionary process begins. Memes are transmitted by copying from person to person, compete to be copied and selected, and thereby evolve. This suggests the perhaps surprising implication that once imitation occurs (whether in human or non-human animals or human-made meme machines), language may spontaneously appear through the competition between sounds to be copied (Blackmore, 1999).

There is evidence from both computer simulations and studies of robots to confirm this. For example, Luc Steels (2000), a computer scientist at the Free University of Brussels, has built the 'talking heads': robots that can make sounds, detect each other's sounds, and imitate them. They have simple vision and categorisation systems, and can track each other's gaze while looking at scenes including coloured shapes and objects. By imitating each other when looking at the same thing, they develop a lexicon of sounds that refer to the shapes they are looking at, although a listening human may or may not understand their words.

Developing grammar proved harder, but a breakthrough occurred when Steels realised that the speaker could apply its language comprehension system to its own utterances, either before speaking them or after a failure in communication.

'How could a slow, mindless process build a thing that could build a thing that a slow mindless process couldn't build on its own?'

(Dennett, 2017, p. 77)



FIGURE 12.16 • The 'talking heads' are robots that imitate each other's sounds while looking at the same object. From this interaction, words and meanings spontaneously emerge. Could human language have emerged the same way? Does the robots' use of meaningful sounds imply consciousness?

'Robots that imitated humans would acquire an illusion of self and consciousness just as we do'

(Blackmore, 2003, p. 19)

This required a re-entrant mapping in which the output from speech production was internally streamed as input to understanding. Steels (2003) argues that this is not only comparable with re-entrant systems in the human brain but also explains why we have such persistent inner voices chattering away to ourselves. This 'inner voice', he suggests, contributes to our self-model and is part of our conscious experience.

Would imitating robots, or artificial meme machines, then invent self-reference, with words for 'I', 'me', and 'mine'? If so, a centre of narrative gravity would form (Dennett, 1991), and the machines would become deluded into thinking they were an experiencing self. Similarly, the memes they copied might gain a replication advantage by being associated with the words 'I',

'me', and 'mine', and so a selfplex would form, with beliefs, opinions, desires, and possessions, all attributed to a non-existent inner self.

This approach implies that machines capable of imitation would be qualitatively different from all other machines, in the same way that humans differ from most other biological species. Not only would they be capable of language, but their ability to imitate would set off a new evolutionary process – a new machine culture. Early research with groups of very simple imitating robots is already exploring the emergence of artificial culture in robot societies (Winfield and Griffiths, 2010). One question for the future would be whether we and the new imitation machines would share a common expanded culture or whether they would imitate in ways that we could not follow. Either way, they would be conscious for the same reason we are: because they have constructed a false notion of self as the subject experiencing a stream of consciousness. They would become deluded machines believing there was something it's like to be them.

I'M SURE IT LIKES ME

When Tamagotchis hit the playgrounds in the mid-1990s, children all over the world starting caring for mindless little virtual animals, portrayed on tiny, low-resolution screens in little hand-held plastic boxes. These young carers took time to 'wash' and 'feed' their virtual pets, and cried when they 'died'. Soon the craze was over. The Tamagotchi meme had thrived on children's caring natures, but then largely fizzled out, perhaps because the target hosts quickly became immune to such a simple trick. More recently, people have got just as hooked on using their phones to find and fight battles with 3D animals lurking in real environments, with stories of players falling off cliffs and wandering into former concentration camps in search of the Pokémon GO creatures.

We humans seem to adopt the intentional stance towards other people, animals, toys, machines, and digital entities on the flimsiest of pretexts. This tactic of attributing mental states to other systems can be valuable for understanding or interacting appropriately with them, but is not an accurate guide to how those

other systems really work. For example, consider the wall-following robots whose useful behaviour emerged from a couple of sensors and some inherent bias. Or consider the equally simple robots that can gather pucks into heaps. They roam around with a shovel-like collector on the front which either scoops up any pucks they bump into or drops them when it has too many. In consequence, after some time, the pucks are all collected into piles. Observers readily assume that the robots are ‘trying’ to gather up the pucks. In reality, the robots have no goals, no plans, no knowledge of when they have succeeded, and no internal representations of anything at all.

This should remind us that our attributions of intentionality are not to be trusted. A strong impression that a given machine is trying to achieve a goal is no guarantee that it is. And perhaps the same logic should apply when thinking about people as about other machines. As Brooks puts it ‘we, all of us, overanthropomorphize humans, who are after all mere machines’ (2002, p. 175).

The intentional stance is the attribution of beliefs to a rational agent, and we adopt it all the time. We may be less willing to adopt the ‘phenomenal stance’, by attributing full subjectivity (including consciousness and emotions) to others (Metzinger, 1995b; Robbins and Jack, 2006). Yet we may feel sorry for cartoon characters, love and cherish our dolls and teddies and even our cars, and cringe when we accidentally step on a worm. If asked whether we truly believe that Mickey Mouse, our favourite dolls, or ants and woodlice have subjective experiences, we may emphatically say ‘no’, and yet still behave towards them as though they do. In this way, our natural tendencies to treat others as intentional, sociable, and feeling creatures all confuse the question of artificial consciousness.

This confusion is likely to get deeper as more and more interesting machines are constructed. Among those already with us are some specifically designed to elicit social behaviour from the people they meet. One of Cog’s designers (Concept 12.2), Cynthia Breazeal, was once videotaped playing with Cog. She shook a whiteboard eraser in front of Cog; Cog reached out and touched the eraser; Cynthia shook it again. It looked to observers as though Cynthia and Cog were taking turns in a game.

In fact, Cog was not capable of taking turns; that was a skill scheduled for years further on in its developmental chart. It seemed that Breazeal’s own behaviour was coaxing more abilities out of Cog than had been put in. This set her thinking about how humans interact socially with machines, and to find out more she built Kismet (Breazeal, 2001), a human-like head with some simple abilities built in, and one of the first and best-known ‘social robots’. Many people behaved as though Kismet were alive. They behaved as though Kismet were conscious. More recently, the ‘human-like social robot’ Brian 2 has been designed to be capable of emotional body language, using a variety of postures and movements identified in human interactions (McColl and Nejat, 2014). The robotic head EMYS is part of a project to create robotic companions for humans. It has a roughly spherical head made of three moveable discs for conveying basic emotions like anger, disgust, sadness, and surprise. Thirty-three percent of the 8–12-year-old children surveyed thought it had emotions, and rated it as having a very positive personality on the ‘big five’ personality factors (Kędzierski et al., 2013). These developments show how readily people infer consciousness in non-human machines, confirming just how easy it is to take the intentional stance.

‘We behavioral and cognitive scientists have been trained to view anthropomorphism as a seductive demon at which to shake objectivistic garlic’

(Reber, 2016, p. 3)




FIGURE 12.17 • A variety of social robot heads (Kędzierski et al., 2013).




FIGURE 12.18 • Cynthia Breazeal with Kismet, the sociable robot. Kismet had four colour video cameras, an auditory system, motors with 15 degrees of freedom controlling face movements, and a vocalisation system able to communicate personality and emotional quality.

Other social robots have been developed to study human-robot interactions and bonding, as well as for potential commercial exploitation. For example, iCat is a small desktop cat-like robot able to play games such as tic-tac-toe (noughts and crosses) and to assume different personality traits. aMuu is an emotional robot designed to explore the possibilities for future home robots; Probo is an elephant-like head and torso that is soft enough to hug; and KASPAR is a child-sized humanoid robot with movable head, arms, hands, and neck, and silicon-rubber mask face, designed to study gestures, facial expressions, synchronisation with human behaviours, and imitation. People happily touch, talk, and play games with these robots, bringing emotionally embodied responsiveness into a realm where it has often been thought lacking (Stuart, 2011). Indeed, elderly people with mild Alzheimer’s disease seem to use more gestures and physical contact with a teleoperated robot called Telenoid than they do with a human carer (Kuwamura, Nishio, and Sato, 2016). None of these robots looks remotely convincing as a human, and so they avoid the uncomfortable reaction that comes with a robot that is somehow too uncannily close for comfort.

You might jump to the obvious conclusion that the human provides all the real meaning and the only source of consciousness in these interactions. You might be confident that iCat, Probo, and KASPAR no more have relationships with you than the cats and dogs some people let themselves believe care about them. You might be sure these robots cannot be conscious because they’re just piles of metal and fabrics with a set of simple routines, just as you might be confident that the fish on your plate could not have been conscious because it had the wrong neural architecture. But it is worth pausing first to note some similarities between us and the social robots.

Brooks says of Kismet, ‘There was no place that everything came together and no place from which everything was disseminated for action’ (2002, p. 92). In

‘we, all of us, overanthropomorphize humans, who are after all mere machines’

(Brooks, 2002, p. 175)

other words, Kismet has no Cartesian Theatre – but then, as we concluded in Chapter 4, we don't either. Like Kismet, we humans have a subsumption architecture. That is, evolution kept whatever worked, dropped what did not, and piled new routines on top of old ones in haphazard interacting layers without an overall plan. This is often how robots improve as a team of scientists works on them.

All these robots are much simpler than us, but let's imagine some fanciful future descendants of today's social robots who are even more skilful. Imagine CREEPI (Conscious Robot with Evolved Emotional and Phenomenal Intelligence), who is still just a mass of metal limbs, motors, and chips, but has soft human-like skin that can convey subtle facial expressions, and eyes that cry wet tears accompanied by convincing sobs, activated by systems that respond to the person who is in front of it. Imagine that CREEPI can respond to emotions displayed by a human: laughing when the human laughs, or sympathising and comforting someone who appears upset. Imagine that CREEPI is even more sensitive to other people's emotions than most humans. What would you say now? Would you still be sure that CREEPI is just a pile of bits, or would you think that maybe it was conscious?

An obvious response is this. We know that simple systems can mislead us into thinking they have plans, goals, and beliefs when they haven't, and that more complex ones can mislead us even more. So we should not be fooled. We should conclude only that CREEPI acts *as if* it is conscious, when it is not *really* conscious.

An alternative response is this. There is no dividing line between *as-if* and *real* consciousness. Being able to sympathise with others and respond to their emotions is one part of what we mean by consciousness. Today's social robots have a little bit of it and CREEPI has a lot. CREEPI is not conscious in the way that we are because it is a social machine without other abilities, but within its limited domain its consciousness is as real as any. Maybe, indeed, our own complexity is what misleads us into believing that consciousness can increase only with growing complexity.

Which is right? And how can we find out?

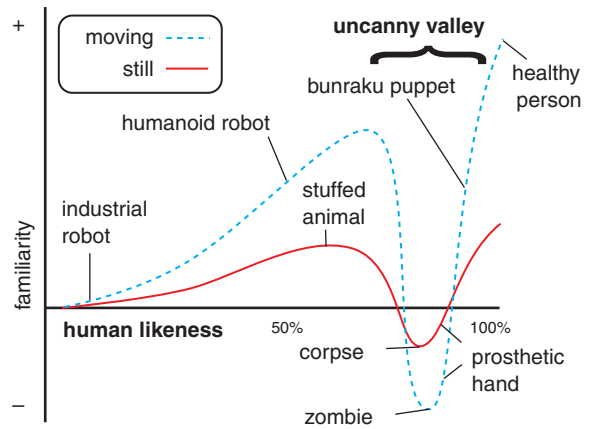


FIGURE 12.19 • The 'uncanny valley' is the name given to the dip in a hypothetical graph plotting emotional reactions against the similarity of an artefact to a human. The most negative response occurs to robots or toys that are 'almost human'. Movement increases the effect.

READING

Chalmers, D. (1996). Strong artificial intelligence. In D. Chalmers, *The conscious mind: In search of a fundamental theory* (pp. 313–322). New York: Oxford University Press.

Defends strong AI against objections, including the Chinese Room.

Dennett, D.C. (1998). The practical requirements for making a conscious robot. In D.C. Dennett, *Brainchildren: Essays on designing minds* (pp. 151–170). London: Penguin.

Dennett's involvement in the Cog project, reflecting on embodiment, symbol grounding, and whether an AI could ever *care* about anything, or be an authority on its own internal states.

Harnad, S. (2007). Can a machine be conscious? How? *Journal of Consciousness Studies*, 10(4–5), 67–75.

How this question relates to the problem of other minds, and how Turing-testing (inferring mental states from behaviours) is our only possible response to both.

Searle, J. (1980). Minds, brains, and programs. *Behavioral and Brain Sciences*, 3, 417–457. (Also reprinted in Hofstadter and Dennett, 1981, with commentary by Hofstadter, pp. 353–382.)


Searle's classic paper on the Chinese room, and its many responses.

Sloman, A. and Chrisley, R. (2003). Virtual machines and consciousness. *Journal of Consciousness Studies*, 10(4–5), 133–172.

Argues that building artificial systems (like the CogAff architecture) can contribute to the study of consciousness, including sections on qualia, zombies, introspection, and evolution.

Turing, A. (1950). Computing machinery and intelligence. *Mind*, 59, 433–460. (Partially reprinted with commentary in Hofstadter and Dennett, 1981.)

A classic on the question 'can machines think?'



Borderlands

SECTION

Five



Altered states of consciousness

CHAPTER

THIRTEEN

One conclusion was forced upon my mind at that time, and my impression of its truth has ever since remained unshaken. It is that our normal waking consciousness, rational consciousness as we call it, is but one special type of consciousness, whilst all about it, parted from it by the filmiest of screens, there lie potential forms of consciousness entirely different. We may go through life without suspecting their existence; but apply the requisite stimulus, and at a touch they are there in all their completeness, definite types of mentality which probably somewhere have their field of application and adaptation. No account of the universe in its totality can be final which leaves these other forms of consciousness quite disregarded. How to regard them is the question.

(James, 1902, p. 388)

'our normal waking consciousness [...] is but one special type of consciousness, whilst all about it, parted from it by the filmiest of screens, there lie potential forms of consciousness entirely different'

(James, 1902, p. 258)

This is, indeed, the question, and in this chapter we will consider the nature of some of these 'other forms of consciousness', including drug-induced states, hypnosis, mental illness, and mindfulness.

DEFINING ASCS

James's 'other forms of consciousness' would now be called 'altered states of consciousness' or ASCs – a concept that seems simple but is notoriously difficult to define. I get drunk and so feel and act differently; I recover from depression and wonder how life could ever have felt so unliveable; I feel like a calmer person on the meditation mat. In all these cases, something has obviously changed, but what? As soon as we start to think more deeply about altered states of consciousness, the problems begin.

Should we define ASCs objectively or subjectively? Taking objective definitions first, we might define ASCs in terms of how they were induced, for example by mind-altering drugs or by hypnosis or progressive relaxation. Then we might label different drug-induced states according to which drug the person took, saying that someone was drunk on alcohol, stoned on cannabis, tripping on LSD, or spaced-out or loved-up on ecstasy. But numerous problems make this unsatisfactory. How do we know whether your trip is similar to mine? How do we know whether two slightly different drugs produce the same, or different, ASCs? And how can we measure the similarity so as to make such decisions? Then there is dosage. How much cannabis does someone need to take to say that they have reached the stoned ASC? Even if two people do experience similar states, the dose required may be quite different for each person. We might define a hypnotic state in terms of a particular hypnotic procedure, but that same procedure might have no effect on someone else. Defining and categorising ASCs by the way they are induced is not satisfactory.

Perhaps a better solution is to define ASCs on the basis of physiological and behavioural measurements, such as heart rate, cortical oxygen consumption, ability to walk in a straight line, or expressions of emotion. One problem here is that very few ASCs are associated with unique physiological patterns (a partial exception is sleep, Chapter 15) or with physiological or behavioural changes that map directly onto changes in experience. As methods improve, we may find consistent patterns, enabling us to define the states in terms of those measures. But it may turn out that very small changes in physiology can be associated with large changes in subjective state, and vice versa, so that no direct mapping is possible. For the moment, we should be careful about defining a state of consciousness (SoC) in terms of physiology. There is a danger of losing the very essence of ASCs, which is how they feel for the person concerned.

The alternative is to define ASCs subjectively, and although objective definitions are sometimes used (e.g. Revonsuo et al., 2009), this is the most common strategy. The term ASC was first formally defined by American psychologist Charles Tart as 'a qualitative alteration in the overall pattern of mental functioning, such that the experiencer feels his consciousness is radically different from the way it functions ordinarily' (1972a, p. 1203). An early textbook on consciousness described the ASC as 'a temporary change in the overall pattern of subjective experience, such that the individual believes that his or her mental functioning is distinctly different from certain general norms for his or her normal waking

state of consciousness' (Farthing, 1992, p. 205). Similar definitions persist in psychology textbooks. One popular volume says that 'an *altered state of consciousness* exists whenever there is a change from an ordinary pattern of mental functioning to a state that seems different to the person experiencing the change' (Nolen-Hoeksema et al., 2014, p. 640).

Such definitions capture the basic idea of ASCs but raise problems of their own. First, they compare ASCs with a normal SoC, but what is normal? Normality for one person might range through bleary-eyed breakfast-eating to concentrated hard work, and from relaxing alone with music to having coffee with friends. Arguably the 'breakfast-eating state of consciousness' differs as much from 'having coffee with friends' as being stoned differs from being straight, and yet most people would unhesitatingly agree which were 'normal'. So the subjective definition of ASCs depends on comparing them with normal states, but we cannot pin those down either.



FIGURE 13.1 • The experience might be amazing, but the words never seem to do it justice.



PRACTICE 13.1 IS THIS MY NORMAL STATE OF CONSCIOUSNESS?

As many times as you can, every day, ask yourself '***Is this my normal state of consciousness?***' When you have decided, you might like to ask some other questions. How did you decide? What is normal about it? Is it always obvious what state you are in, and if so, why? If not, what does this tell you about ASCs?

Another problem is inherent in the whole idea of subjective definitions: they may help us to decide for *ourselves* whether we are in an ASC, but as soon as we try to tell others, our words become objective behaviour from their point of view. Also, think of the drunk who staggers about claiming that he feels perfectly normal, or the first-time marijuana smoker who giggles at her own hand for ten minutes while insisting that the drug has had no effect. In these cases, we may think that physiological measures would be more appropriate than words. And even when the person's own words seem to be the best measure, there are still problems because ASCs are notoriously hard to describe, and different people have different prior experiences, different expectations, and different ways of describing things. Training may help, but this raises other problems such as how to compare the experiences of trained explorers with those of novices.

You may have noticed that lurking among these problems is an old familiar one. Is there really such a thing as a conscious experience that exists apart from the things people do and say about it? Or is consciousness itself nothing more than those behaviours and descriptions, as claimed, for example, by eliminative

materialists, identity theorists, and some functionalists? If such theories are correct, then we should be able to understand ASCs fully by studying the physiological effects and behaviour, and there should be no mystery left over. Yet for many people this does not do justice to what they feel. They enter a deep ASC and everything seems different. They struggle to describe it, but somehow the words are not enough. They *know* what they are experiencing but cannot convey it to anyone else. They *know* that their conscious experience has changed in ways their behaviour and words cannot convey. Are they right?

WHAT IS ALTERED IN AN ASC?

‘What is altered in an altered state of consciousness?’ is a strange but interesting question. Optimistically, we might say that ‘consciousness’ has changed. If this is so, studying what is altered should reveal what consciousness itself really is. Sadly, everything we have learned so far shows how difficult this is. We do not know how to measure changes in something called consciousness in isolation from changes in perception, memory, or other cognitive-emotional functions, so to study ASCs we must start by studying how these functions have changed.

All the definitions given above, as well as comparing ASCs with a normal state, mention a change to ‘mental functioning’. So which kinds of functioning are involved?

Farthing (1992) provides a list: 1) attention, 2) perception, 3) imagery and fantasy, 4) inner speech, 5) memory, 6) higher-level thought processes, 7) meaning and significance, 8) time perception, 9) emotional feeling and expression, 10) arousal, 11) self-control, 12) suggestibility, 13) body image, and 14) sense of personal identity. In one way or another, this list probably covers all mental functions, suggesting that ASCs cannot be fully understood without understanding changes to the whole system. Some ASCs involve changes to all these functions, while others primarily involve just one or two, and we shall meet many examples of these in the rest of this chapter. For now, we might pick out just three major variables that often change during ASCs: attention, memory, and arousal.

Attention can change along two main dimensions: direction and focus. First, attention may be directed ‘inwards’ or ‘outwards’. For example, in daydreaming sensory input is largely ignored and attention is focused on trains of thought and imagery. Good hypnotic subjects may ignore the world around them and concentrate entirely on the hypnotist’s suggested fantasies. Many methods for inducing ASCs manipulate this dimension either by reducing sensory input, as in meditation or deep relaxation, or by overloading it as in some ritual practices. Second, attention may be broadly or narrowly focused. Someone high on marijuana may attend finely to the leaf pattern on the carpet for many minutes at a time. Such a change in attention can seem to profoundly affect subjective states, but the effects cannot be cleanly separated from the associated changes in perception, memory, and emotion. For example, the leaf pattern might look quite different from normal, become of overwhelming significance, bring up long-lost childhood memories, and raise deep emotions – or gales of laughter.

Second, memory changes occur in many ASCs and are linked with effects on thinking and emotion. For example, many mind-altering drugs reduce short-term memory span. This has a debilitating effect on conversation if you cannot

remember what you started out to say before you finish the sentence, but it can also create more focused attention on the here-and-now, and even a sense of liberation. Time can seem to speed up, slow down, or change completely, an effect that has long been linked with changes in memory. For example, a doctor experimenting with cannabis more than a century ago noted many effects including a dry mouth, aimless wandering, slurred speech, freedom from worry, and an irresistible tendency to laugh. For him,

The most peculiar effect was a complete loss of time-relation; time seemed to have no existence. I was continually taking out my watch, thinking that hours must have passed, whereas only a few minutes had elapsed. This, I believe, was due to a complete loss of memory for recent events.

(Dunbar, 1905, p. 68)

The third general variable is arousal. Some states of meditation are characterised by very low arousal and deep relaxation (Holmes, 1987), and more drastic practices can reduce the metabolic rate so far that little food and oxygen are required. In such a state, trained yogins may stay immobile for long periods and may even be buried alive for days at a time, even though most of us would die in the same circumstances. However, learning to meditate requires great mental effort, and with some methods arousal is increased rather than reduced (Lumma, Koko, and Singer, 2015). At the other extreme are ASCs of the highest arousal, such as religious and ritual frenzies, or speeding on amphetamines. Changes in arousal can affect every aspect of mental functioning.

Thinking about these three variables, we might imagine some kind of three-dimensional space in which different ASCs can be positioned – or, more realistically, a very complex multidimensional space within which all possible ASCs might be found: a phenomenal state space, or phenospace (Metzinger, 2009). If states of consciousness (SoCs) could be accurately mapped within such a space, we might understand how each relates to the others, how each can be induced, and how to move from one state to another. But although many attempts have been made, the task is not easy.

MAPPING STATES OF CONSCIOUSNESS

Imagine a vast multidimensional space in which a person's current state is defined by hundreds or even thousands of variables. This is just too confusing to work with. To make the task more manageable, we need to answer two main questions: first, can we simplify the space and use just a few dimensions, and if so, which ones; second, how discrete are the individual SoCs? Is it possible to occupy any position in the multidimensional space, or are possible SoCs separated from each other by impossible areas?

The early psychophysicists tried to map visual and auditory sensations in multidimensional spaces, but the first attempt at systematic mapping of states of consciousness was made by Tart (1975). He described a simple space with




FIGURE 13.2 • It is easy to imagine altered states connected up in a vast space, but difficult to turn this idea into a realistic working map. How discrete are the different states? Where do the paths between them go? How many dimensions are there, and how many would we need to use to make an effective map?

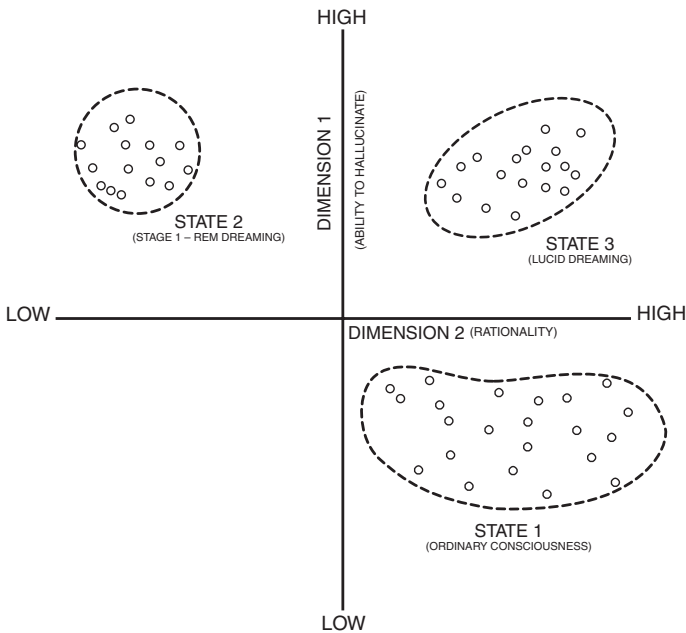


FIGURE 13.3 • Tart's plot of three discrete altered states in a space with two dimensions: irrationality and ability to hallucinate (after Tart, 1975, Fig. 5.1).

two dimensions: irrationality and ability to hallucinate. By plotting a person's position in this space, he imagined just three major clusters corresponding to the states of dreaming, lucid dreaming (Chapter 15), and ordinary consciousness. All other positions in the space cannot be occupied, or are unstable. So you may briefly hover between waking and dreaming, but this state is unstable and rapidly gives way to one of the others. For this reason, Tart refers to the occupied areas as 'discrete states of consciousness'. To move out of such a region, you have to cross a 'forbidden zone' where you cannot stably function or have experiences, until you reach a discretely different experiential space. In other words, you can be here or there, but not in between. Just how many states are discrete like this we do not know: Tart's scheme was only a limited and quite informal way of starting to map states into a space.

A second and more systematic two-dimensional space is described by Steven Laureys (2005). His dimensions are completely different: level of arousal and awareness of environment and self. *Arousal* refers to physiological wakefulness or the 'level' of consciousness and is dependent on the brainstem arousal system. *Awareness of environment and self* refers to the 'content' of consciousness and requires a functionally integrated cortex with its subcortical loops. A simple diagram shows that for most states level and content are positively correlated. As Laureys puts it, 'You need to be awake in order to be aware (REM sleep being a notable exception)' (2009, p. 58). Other exceptions discussed by Laureys are the vegetative state, sleepwalking, and some kinds of seizure, all of which involve some wakefulness with no apparent awareness.

The AIM model is a three-dimensional map developed by American psychiatrist and sleep researcher Allan Hobson (1999), and named after its dimensions. *Activation energy* is similar to arousal and can be measured, for example, using EEG. *Input source* can vary between entirely 'external' or entirely 'internal' sources of information. *Mode* is the ratio of amines to cholines. During waking, the amine neurotransmitters and neuromodulators, including noradrenaline and serotonin, predominate and are essential for rational thought, volition, and directing attention. During REM sleep acetylcholine takes over and thinking becomes delusional, irrational, and unreflective. The ratio of these two is Hobson's *mode*.

States can now be positioned in what Hobson calls 'brain-mind space' by measuring them along these three dimensions. He stresses that it is an entirely artificial model, yet is based on specific data and recognises the continuously changing nature of brain-mind states. Unlike in Tart's early model, any state of the brain-mind can be positioned within it, and any area in the space can in theory be occupied.

Other dimensional models derive from studies of specific corners of the vast consciousness state-space. A survey of 'meditation depth' amongst 300 yoga, Buddhist, and TM meditators resulted in another three dimensions: mystical

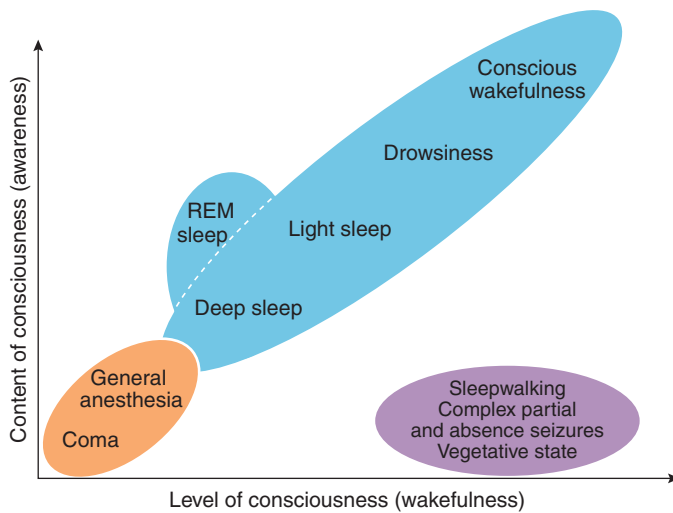


FIGURE 13.4 • Oversimplified illustration of the two major components of consciousness: the level of consciousness (i.e. wakefulness or arousal) and the content of consciousness (i.e. awareness or experience). In normal physiological states (blue) level and content are positively correlated (with the exception of dream activity during REM sleep). Patients in pathological or pharmacological coma (that is, general anaesthesia) are unconscious because they cannot be awakened (pink). Dissociated states of consciousness (i.e. patients being seemingly awake but lacking any behavioural evidence of ‘voluntary’ or ‘willed’ behaviour), such as the vegetative state or much more transient equivalents such as absence and complex partial seizures and sleepwalking (purple), offer a unique opportunity to study the neural correlates of awareness (Laureys, 2005, p. 556).




FIGURE 13.5 • Hobson’s AIM model describes ‘brain-mind space’ using three dimensions: A for activation energy (low to high), I for information source (internal to external), and M for mode (cholinergic to aminergic). States of consciousness can be positioned in the space using data from behavioural and physiological studies. (See Figure 15.3 for a more detailed version of the model.)

experience (bliss, contact with a higher force), nirvana (absence of thought, total absorption), and mental and bodily relaxation (reduction of tension) (Ott, 2001). Other maps made by spiritual practitioners range from one-dimensional evolutionary types (Chapter 10) to more complex spaces including hard-to-define concepts like the cosmos, ego, faith, karma, energy, and the unconscious.



PROFILE 13.1

Thomas Metzinger (b. 1958)



Thomas Metzinger is no one. A German philosopher, and director of research groups in neuroethics and theoretical philosophy at the Johannes Gutenberg University of Mainz, he claims that no one ever had or was a self.

In *Being No One*, and his later book *The Ego Tunnel*, he argues that what we take to be persistent entities are really ongoing processes: the contents of transparent self models. He has experienced and written about many 'altered states', including lucid dreams, out-of-body experiences, meditation, and drug-induced experiences, and is concerned about the ethical implications of our rapidly advancing phenotechnology. When we can choose which areas of phenospace we want to visit, or can enhance our cognitive skills with specially tailored chemicals, we will need to take responsibility for the consequences. Hence much of his work is in the new field of neuroethics. Metzinger was a co-founder of the Association for the Scientific Study of Consciousness and has edited books on conscious experience and the NCCs.

A psychological and neurobiological review of ASCs (Vait et al., 2005) included states experienced spontaneously, stimulated by physical or psychological means, or caused by illness, resulting in a four-dimensional model. The dimensions are *activation* (low to high arousal), *awareness span* (a narrow to broad range of 'contents available to attention and conscious processing', p. 114), *self-awareness* (diminished to heightened), and *sensory dynamics* (reduced to heightened sensation). The authors present their four dimensions as a first step towards constructing what they call the 'C-space': the space of states of consciousness. The counterpart to the C-space is the 'B-space': the space of functional brain states. The challenge is to create mappings between the two, whether these are understood as strict one-to-one mappings or as one-to-many or many-to-many mappings. In any case, they argue that with state-space approaches we should only ever expect the locations in both spaces to be 'blurred' (determined with limited resolution), meaning the final mappings will always be coarse-grained and probabilistic (p. 119).

This really is a big challenge – to bring together subjective experiences of altered states with what we know of the underlying physiology. Some people have an enormously wide range of experiences over their lifetimes. They may have taken many different drugs, practised meditation, yoga, or Tai Chi over long periods, and used TMS, sensory isolation tanks, or other kinds of 'phenotechnology', as Thomas Metzinger calls it. In these ways, they may have gained a good personal understanding of how to shift from state to state, how to maintain or leave a given state, and how each state can be used, but as yet this kind of understanding has not been systematically integrated into academic research.

We may imagine a future in which we thoroughly understand the nature of phenospace and the various technologies that can move us around within it. Metzinger claims that with this knowledge we should in principle be able to design our own ego tunnels (conscious experiences which we attribute to our selves) by tinkering with the hardware that supports them.

Whether the desired phenomenal content is religious awe, an ineffable sense of sacredness, the taste of cinnamon, or a special kind of sexual arousal does not really matter. So, what is your favourite region of phenospace? What conscious experience would you like to order up?

(2009, pp. 220–221)




FIGURE 13.6 • The chemical structures of some well-known psychoactive drugs. Many of these resemble the structure of neurotransmitters.

Although we are a very long way from that depth of understanding, our phenotechnology is rapidly improving and will bring with it ethical and political consequences. These concern both the individual whose consciousness is altered and societies that have to decide whether any technologies or areas of phenospace should be made illegal, and if so, how to enforce such prohibition.

Many procedures can be used to explore the far reaches of phenospace, and some life experiences bring us closer to them whether we want it or not. In Chapter 15, we will explore spontaneously occurring 'altered states' like dreaming and out-of-body experiences. Here we consider mind-altering drugs (a chemical trigger), meditation and hypnosis (psychological routes), and mental illness (pathological causes).

DRUG-INDUCED STATES

Psychoactive drugs are all those that have effects on mental functioning or consciousness (see the website for more on the main categories and their mechanisms). They are found in every society, and human beings seem to have a natural appetite for taking them (Weil, 1998). They all work by changing the action of endogenous neurotransmitters or neuromodulators. For example, they may increase a neurotransmitter's effect by mimicking it, stimulating its release, or blocking its reuptake so that its effects last longer, or may reduce the effects by inhibiting release or blocking its reception in the post-synaptic membrane. One reason the mind-altering effects of drugs can be so dramatically wide-ranging is that even a single neurotransmitter can be active in many different regions of the brain. By knowing the mode of action of a drug and understanding the system it affects, we should in principle be able to understand precisely why each drug has the effect it does.

Psychoactive drugs can be broadly classified into several major groups. All have distinct effects on the brain and on experience (Pace-Schott and Hobson, 2007). Discussed in more detail here are a few from different groups that have the most profound effects on consciousness.

STIMULANTS

Many designer drugs are related to amphetamine, perhaps the best known being 3,4-methylenedioxy methamphetamine, MDMA, or ecstasy. MDMA has three main effects in the brain: inhibiting serotonin reuptake, and inducing the release of serotonin and dopamine. Serotonin plays a major role in regulating mood and sleep, and dopamine helps mediate reward-motivated behaviour as well as interpretive responses to self, other, and environment. So, not surprisingly, MDMA has a mixture of amphetamine-like and psychedelic effects, including increased energy, enhancement of tactile and other sensations, and feelings of love and empathy, for which it is sometimes referred to as an 'empathogen' or 'entactogen' (Holland, 2001; Saunders, 1993). The effects, as with so many other psychoactive drugs, are highly dependent on the setting in which it is taken. At parties and clubs, the increased energy makes dancing all night easy, and bombardment with music and light adds to the effects, but MDMA can also be used to enhance intimacy and sex, or solve personal problems. When taken alone, especially in beautiful surroundings such as mountains or the ocean, MDMA can lead to a

'what is your favourite region of phenospace?'

(Metzinger, 2009, p. 221)

'for thousands of years people of all cultures have used psychoactive substances to induce special states of consciousness'

(Metzinger, 2009, p. 230)

profound sense of union with the universe and love for all creation, and some raves are regarded by those involved as spiritual events (Saunders, Saunders, and Pauli, 2000).

Like many amphetamine derivatives, MDMA produces tolerance and is addictive. There is some evidence of long-term damage to the serotonergic system from even moderate use, although the brain may recover with abstinence and the long-term effects are not yet fully known (Holland, 2001). One famous study published in 2002 by George Ricaurte claimed to have shown neurotoxicity, but was forced to be retracted when it was found that methamphetamine had been used instead of MDMA.

People who use MDMA to explore ASCs or for spiritual purposes tend not to take it frequently or mix it with other drugs and may therefore be less likely to suffer any damage associated with overuse and abuse. And research on MDMA use in therapeutic contexts in fact suggests very promising outcomes for conditions like post-traumatic stress disorder and social anxiety (Metzner and Adamson, 2001; Bouso et al., 2008; Mithoefer et al., 2013; Danforth et al., 2016).

ANAESTHETICS

Most anaesthetics do not produce interesting ASCs and have indeed been designed not to do so. However, some anaesthetic gases and solvents, such as ether, chloroform, and nitrous oxide, can induce quite profound ASCs.

When William James penetrated that filmy screen into another form of consciousness, he had inhaled air mixed with nitrous oxide, a gas first isolated by Sir Humphrey Davy at the Pneumatic Institution, a medical research facility in Bristol. The euphoric effects soon led to its being dubbed 'laughing gas' and used for entertainment at exclusive parties. For the same reason, people now fill balloons with nitrous oxide and breathe it in for a brief and enjoyable high. Its pain-killing effects resulted in its use as an early anaesthetic in dentistry and surgery, but it is now most familiar as the 'gas and air' used for pain relief in childbirth.

Davy bravely experimented with many gases by taking them himself, and breathed his first dose of nitrous oxide on 11 April 1799. He described an immediate thrilling, a pleasure in every limb, and an intensification of both vision and hearing; he became enormously excited, shouting and leaping about the laboratory (Jay, 2009). He lost concern with external things and entered a new world of ideas, theories, and imagined discoveries. On returning to normality, he claimed that 'Nothing exists but thoughts'.

This sounds just like what Tart had in mind when he described different SoCs as having different logics, or involving different ways of seeing the world, leading him to the idea of 'state-specific sciences' (Concept 13.1). It sounds as though Davy had been made into a philosophical idealist; that his intellectual beliefs were different in the ASC. A century later, another explorer of anaesthetics said just the same after inhaling ether mixed with air:

Then it dawned upon me that the only logical position was subjective idealism, and, therefore, *my* experience must be reality. Then by degrees

'the opposites of the world, whose contradictoriness and conflict make all our difficulties and troubles, were melted into unity'

(James, 1902, p. 388)

I began to realise that I was the One, and the universe of which I was the principle was balancing itself into completeness.

(Dunbar, 1905, pp. 73–74)

James described nitrous oxide as stimulating an artificial mystical consciousness, in which ‘depth beyond depth of truth’ is revealed and then fades out when the drug wears off, often leaving only nonsense words behind. Yet the sense of meaning and insight remains, and the insights have been compared with those of Zen (Austin, 2006). There is an experience of reconciliation, said James: a monist rather than dualist insight, ‘as if the opposites of the world, whose contradictoriness and conflict make all our difficulties and troubles, were melted into unity’ (1902, p. 388). As James himself said, the question is how to regard these insights.

Ketamine is a dissociative anaesthetic, although it is rarely used for anaesthesia in humans because it can induce schizotypal symptoms and terrifying nightmares, as well as possible long-term harm (C. Morgan and Curran, 2011). Its main action is as an NMDA antagonist but among other effects are inhibiting the reuptake of serotonin, dopamine, and noradrenaline. Ketamine affects attention, disrupting the deliberate directing of attention rather than the capturing of attention from outside (Fuchs et al., 2015). It also disrupts working memory, episodic memory, and semantic memory, with measurable effects lasting for several days. Nevertheless, there is evidence of therapeutic value for schizophrenia, possibly because it reduces activity in brain areas involved in sensory processing and selective attention (Musso et al., 2011), and for severe depression, where it seems to decrease functional connectivity between networks such as the DMN and affective and cognitive control networks (Scheidegger et al., 2012).

As a recreational drug, ‘K’ or ‘Special K’ is used in sub-anaesthetic doses for its weird psychological effects ranging from peace, euphoria, and vestibular sensations of floating and falling to a dissociated state of derealisation and depersonalisation in which things seem distant, unreal, or inexplicable (Stirling and McCoy, 2010). When injected, the effects begin within a few minutes and last about half an hour; when eaten, the effect is much slower and longer-lasting, with after-effects lasting several hours.



CONCEPT 13 STATE-SPECIFIC SCIENCES

● Charles Tart (1972a) proposed the creation of ‘state-specific sciences’, likening SoCs to paradigms in science. Paradigms are general scientific frameworks within which normal science operates and whose assumptions usually remain unquestioned until there are so many anomalies that a scientific revolution or paradigm shift has to occur, leading to a new paradigm (Kuhn, 1962). Within a paradigm, a certain self-consistent logic applies, certain rules are taken for granted, and all data are interpreted within it. Within a different paradigm, other rules apply.

The same is true, says Tart, of SoCs. They too involve rules and ways of seeing things that are self-consistent but different from those that apply in other SoCs, so different states may need different kinds of science. Research would have to be carried out, and results communicated between people, all in the relevant SoC. This would require highly skilled practitioners able to achieve given states, agree that they had achieved those states, and work within them. They might then investigate any natural phenomena, but the ways they did so, and their findings, would make sense only to people also working within that state.

There are no results from SSSs published explicitly as such, although it is possible that some scientists are doing SSS and communicating with each other in ASCs, without revealing this publicly. Tart (2015) reports, for instance, that some mathematicians may rely on ASCs to do creative maths and comprehend others’ work. There is no doubt that many scientific breakthroughs have been made by people who saw their problems differently in an ASC and then brought that insight back to the normal SoC, but this is only halfway there. In any case, the proposal is interesting because it questions the usual assumption that the ‘normal’ state is the only, or best, state in which science or other research can be done.

In a study comparing frequent, infrequent, and ex-users, two-thirds described the most appealing aspects of ketamine as ‘melting into the surroundings’, ‘visual hallucinations’, ‘out-of-body experiences’, and ‘giggleness’ (Muetzelfeldt et al., 2008). Less appealing were worries about ‘memory loss’, ‘decreased sociability’, and addiction (see also C. Morgan and Curran, 2011). People taking ketamine in experimental settings are more susceptible to the rubber-hand illusion in which you feel a fake hand is your own (H. Morgan et al., 2011; see Chapter 4), and many report bodily distortions: ‘My hands look small, but the fingers are really long’, said one; another remarked, ‘My legs look very big and funny shaped, like another person’s’ (Pomarol-Clotet et al., 2006; see also Curran and C. Morgan, 2000). Sometimes these changes in the body schema progress to illusory movements or to out-of-body feelings (Wilkins, Girard, and Cheyne, 2012). And in high doses, there is the famous ‘K-hole’, an experience of extreme dissociation, derealisation, bodily dissolution, and, in many cases, out-of-body or near-death experiences (Chapter 15). Described as anything from a place of extreme horror to the best time ever, the K-hole is sought by some and feared by others.

On the theory that you can often learn about something by switching it on and off, Richard Gregory (1986) chose an intravenous infusion of ketamine as a way to explore the switching-off of consciousness. Under controlled conditions in the laboratory, he was shown ambiguous figures, random dot stereograms, and words to read, as well as many other tests. The walls began to move; he heard a loud buzzing noise; he felt unreal and floating, as though he were in another world like a bubble full of bright colours and shapes. He even experienced synaesthesia for the only time in his life when he felt the bristles of a brush as orange, green, and red. Interesting as this was, the whole experience was deeply unpleasant for Gregory. He concluded that he had learned little about consciousness and had no enthusiasm for repeating the experience.

Sue’s attempts to induce an out-of-body experience with ketamine were far more pleasant. She had a dose just below anaesthetic level injected intravenously in a pleasant and relaxing environment. ‘I am lying back in some yielding, flowing softness [. . .]. I seem to be disintegrating, falling apart into separate pieces and then into nothing at all. Then back together and flying’ (Blackmore, 1992, p. 273). Despite these interesting sensations, she concluded that it was very different from spontaneous out-of-body experiences (Blackmore, 2017). The physicist Richard Feynman, who experimented with tiny doses of ketamine taken in an isolation tank, reported that it made him feel as though he were an inch to one side, and that with practice he could move down inside or further away from his own body, until ‘everything else was exactly the same as normal, only my ego was sitting outside, “observing” all this’ (Feynman and Leighton, 1985, p. 333).

Ketamine is also used in different settings as a sacred or therapeutic drug. It is then as much a psychedelic as an anaesthetic, used to explore the grand questions of birth, life, and death (Jansen, 2001). Gregory’s unpleasant experience in the laboratory illustrates how important set (or state of mind) and setting (environment) are in establishing the effects of psychoactive drugs.

PSYCHEDELICS

The effects of drugs in this group are so strange and varied that there is no firm agreement even over their name. We have called them *psychedelics*, meaning

mind-manifesting, but other names are often used. *Psychotomimetic* means madness-mimicking, but this is inappropriate because although existing psychosis can be aggravated by some of these drugs, few features of psychosis are mimicked by them. They are also called *hallucinogenic*, although 'true' hallucinations – in which the person thinks their hallucinations are real – are rare (Julien, 2001; Shulgin and Shulgin, 1991; and Chapter 14). Other terms include *psycholytic*, meaning loosening the mind, and *entheogen*, meaning releasing the god within. Cannabis is sometimes referred to as a minor psychedelic or hallucinogen, with the rest being major psychedelics.

Cannabis. The familiar and beautiful plant *Cannabis sativa* has been used medically for nearly 5,000 years, and as a source of tough fibre for clothes and ropes for even longer (Earleywine, 2002). Cannabis contains hundreds of chemical components, including a range of at least eighty-five cannabinoids, the most important of which are cannabiniol (CBN), cannabidiol (CBD), and tetrahydrocannabinol (THC), the main psychoactive constituent. In the nineteenth century, cannabis (also known as marijuana) was widely used as a medicine. Medical use and knowledge were then restricted by over half a century of prohibition (Booth, 2003), but there are some signs of this starting to relax, with the medicinal benefits of cannabis use becoming more widely accepted. Nineteenth-century scientific explorers of cannabis and the artist members of the Club des Hashischins, such as Balzac and Baudelaire, ate hashish. This is a dark brown or reddish solid derived from the resin scraped from the female flowers, leaves, and stems, and sometimes including powdered flowers and leaves. Cannabis can also be made into a tincture with alcohol or a drink mixed with milk, sugar, and spices, or cooked with butter or other fats in chocolate, cakes, or savoury dishes. As a recreational drug in the twenty-first century, it is most often smoked in the form of hash mixed with tobacco or burnt alone in special pipes, as oil smoked in electronic cigarettes, or as grass, the dried leaves and buds smoked on their own or with tobacco or dried herbs.

As with any drug, smoking makes for rapid absorption into the bloodstream by avoiding enzymes in the digestive system that can break down some constituents, and also allows for easy control over the dose. When eaten, the effect is slower and longer-lasting, and control is more difficult. The main active ingredients are all fat-soluble, and some can remain dissolved in body fat for many days or even weeks after smoking. With its complex and varying mixture of psychoactives, cannabis nicely illustrates the difference between natural psychoactive mixtures, which also include ayahuasca and drugs derived from mushrooms and cacti, and the simpler or starker effects of synthetic psychedelics. When one or more of the active ingredients is isolated, the rich and varied psychological effects are usually lost (Weil, 1998).



FIGURE 13.7 • *Cannabis sativa* is a beautiful fast-growing annual that thrives in a wide variety of climates, shown here ready for harvest in a greenhouse in Britain. The leaves and flowering heads are smoked as grass, and hash can be made from the resin.

**IS THIS MY
NORMAL STATE OF
CONSCIOUSNESS?**

Describing the subjective effects of cannabis is not easy, partly because ‘Most people cannot find the words to explain their sensations’ (Earleywine, 2002, p. 98) and partly because the effects differ so widely from person to person. Some people become self-conscious, disorientated, and paranoid and are disinclined to repeat their experience, while others experience delight, novelty, insight, or just relaxation and go on to strike up a positive, sometimes lifelong, relationship with the drug (Sagan, 1971). Nevertheless, research has revealed some typical effects. In the first major survey of cannabis use, Tart (1971) asked over 200 questions of 150 people, mostly Californian students, who had used the drug at least a dozen times. Other studies have subsequently administered cannabis, or just THC, in the laboratory and recorded the effects.

Users report many emotional effects, including euphoria and relaxation at lower doses, and fear and paranoia at higher doses. Sensory effects include enhancement of all the senses, enhanced depth perception, increased sexual responsiveness and enjoyment, slowing of time, widening of space, and a focus on the present. Synaesthesia is sometimes reported at high doses. Openness to experience increases and some people find a sense of the sacred or divine. Memory, especially short-term memory, is often felt to be impaired. Creative thought and personal insight are often reported, but so are mental fogginess, slowed thinking, and inability to read.

Laboratory studies show that the perceived effects on memory are roughly accurate, with short-term memory severely disrupted while episodic and semantic memory remain generally good. On the other hand, the enjoyable experience of enhancement of the senses is not supported by objective tests: ‘people think marijuana can enhance some visual processes, and laboratory research suggests it actually impairs some of them’ (Earleywine, 2002, p. 105). Planning, problem-solving, decision-making, and basic motor coordination are also affected (Crean et al., 2011). There is also some evidence for longer-term effects on brain development (especially neural connectivity in particular areas such as the precuneus and fimbria) and on other aspects of physical and mental health (Volkow et al., 2014), but this evidence is almost exclusively correlational rather than demonstrating causation. Varieties of cannabis vary in ingredients and effects. For example, cannabidiol (CBD) is found at high levels in natural cannabis but much lower levels in modern skunk varieties bred to fetch high prices on the illegal market. Some studies suggest that CBD may protect against some of the harmful effects of THC (Niesink and van Laar, 2013).

She was sorry, and rather revolted at his dirty hands, but she laughed in a well-bred way, as though it were nothing unusual to her to watch a man walking in a slow dream. Often people display a curious respect for a man drunk, rather like the respect of simple races for the insane. Respect rather than fear. There is something awe-inspiring in one who has lost all inhibitions, who will do anything. Of course we make him pay afterward for his moment of superiority, his moment of impressiveness.

What can we learn about ASCs from this extraordinary mixture of complex effects? Here we have a range of states, sought out by millions of people worldwide, which we seem able to describe only as a mishmash of effects on thought, perception, emotion, and other cognitive functions. As for mapping them, the task seems daunting. Not only is it difficult to position being stoned in relation to other ASCs, but being stoned itself varies widely. Experienced users can readily discriminate between cannabis that induces heady or intellectual experiences, or mellow and relaxing ones, and 'laughing grass' that makes everything funny. They can express clear changes in the nature of their experiences which may or may not be in agreement with results of scientific tests that use measures other than verbal self-report. We are very far from a science of ASCs that can make sense of all this.

He wheeled off his bicycle, feeling Nicole's eyes following him, feeling her helpless first love, feeling it twist around inside him. He went three hundred yards up the slope to the other hotel, he engaged a room and found himself washing without a memory of the intervening ten minutes, only a sort of drunken flush pierced with voices, unimportant voices that did not know how much he was loved.

(F. Scott Fitzgerald, *Tender Is the Night*, 1934)

Major psychedelics

I was [. . .] back in a world where everything shone with the Inner Light, and was infinite in its significance. The legs, for example, of that chair – how miraculous their tubularity, how supernatural their polished smoothness! I spent several minutes – or was it several centuries? – not merely gazing at those bamboo legs, but actually *being* them – or rather being myself in them; or, to be still more accurate [. . .] being my Not-self in the Not-self which was the chair. [. . .] four bamboo chair legs in the middle of a room. Like Wordsworth's daffodils, they brought all manner of wealth – the gift, beyond price, of a new direct insight into the very Nature of Things.

(Huxley, 1954, pp. 20–21, 25)

This is how Aldous Huxley, novelist and author of *Brave New World*, described some of what happened when 'one bright May morning, I swallowed four-tenths of a gramme of mescaline dissolved in half a glass of water and sat down to



ACTIVITY 13.1

Discussing ASCs

People who have experienced ASCs often enjoy talking about them, whether to share their insights, laugh about their exploits, or explore their fears. This needs a supportive and safe environment and you, as leader of any discussion, must decide whether you can provide it or not. In many European countries and now also many states of the USA, cannabis has been decriminalised for medical and/or recreational use, and in some of these areas many other recreational drugs are tolerated, but elsewhere, including at federal level in the USA, anti-drugs laws are stringent. If you cannot talk freely, restrict the discussion to alcohol, sleep, and spontaneous ASCs. You might ask:

Why do you induce ASCs? What do you gain from them?

How can you tell when you have entered an ASC?

Is one person's ASC (such as being drunk, or stoned) the same as someone else's?

Another exercise requires advance preparation but avoids problems of prohibition. Ask participants to bring along a short description of an ASC. This can either be someone else's, for example from a book or website, or their own. They read this out and ask everyone else to guess which ASC is referred to. Discussing how they decided leads naturally to all the other interesting questions about ASCs.

'The legs [. . .] of that chair – how miraculous their tubularity, how supernatural their polished smoothness!'

(Huxley, 1954, pp. 20–21)

'If the doors of perception were cleansed, everything will appear to man as it is, infinite'

(William Blake, *The Marriage of Heaven and Hell*, 1790)

wait for the results' (1954, p. 13). In *The Doors of Perception*, he describes how a vase of three ill-assorted flowers became the miracle of creation; how time and space became insignificant; how his own body seemed perfectly capable of acting without him; and how everything simply *was*, in its own isness or suchness. From these profound experiences, he surmised that the brain works as a reducing valve, preventing our connection with reality, and that drugs can open the valve.

Mescaline, or trimethoxyphenylethylamine, is the main active ingredient in the San Pedro cactus *Trichocereus pachanoi* and in peyote, a small, spineless desert cactus *Lophophora williamsii* that has apparently been used for ritual purposes for at least 7,000 years (Devereux, 1997). Traditionally the top of the peyote is dried to make mescal buttons, which are then chewed to invoke deities and open up other worlds. Mescaline is also produced synthetically and is then used on its own without the complexity of the thirty or so other alkaloids that are found in peyote. Mescaline makes the world seem fantastic and colourful, which is reflected in the art it has inspired, and contributes to 'the conviction that this is a view of the essential nature of the universe' (Perry, 2002, p. 212). This is probably its most characteristic effect, and we will learn more about it in the next chapter. Some users describe mescaline as more hallucinogenic and less self-revealing, or self-destroying, than some of the other psychedelics, especially LSD.

Psilocybin is found in many mushrooms of the *Psilocybe* genus, often called magic mushrooms or sacred mushrooms. They include *Psilocybe cubensis* and *Psilocybe mexicana*, which can (with difficulty) be cultivated, and many other species native to different parts of the world. When it was readily available and legal in the 1960s, Timothy Leary and other members of the 'Harvard psilocybin project' used psilocybin to encourage people to 'turn on, tune in, drop out' (Stevens, 1987).

'Turn on' meant to [. . .] [b]ecome sensitive to the many and various levels of consciousness and the specific triggers that engage them. [. . .]. 'Tune in' meant interact harmoniously with the world around you [. . .]. 'Drop out' meant self-reliance, a discovery of one's singularity, a commitment to mobility, choice, and change.

(Leary, 1983, p. 253)

Psilocybin's effects typically last for 3–4 hours, making it a much more manageable drug than LSD, and for this reason it is preferred for scientific research, including studies on psychedelic-assisted therapy for treating depression (Carhart-Harris et al., 2016a) and for helping reduce anxiety and improve quality of life in the terminally ill (Ross et al., 2016). The risk of adverse psychological effects for healthy users taking ordinary doses is low (Studerus et al., 2011), and the drug is often also claimed to induce mystical and religious experiences. Psilocybin often induces a feeling that one's self is disintegrating and merging with the rest of the world, and this 'ego dissolution' has been found to be associated with decreased functional connectivity between the medial temporal lobe and high-level cortical regions, as well as with reduced interhemispheric communication (Lebedev et al., 2015). Given that psilocybin is meant to be 'mind-expanding', researchers were surprised to find that when profound experiences were induced in an fMRI scanner, only decreases in cerebral blood flow were found, especially in the thalamus and cingulate cortex (Carhart-Harris et al., 2012). Some have taken this to support Huxley's idea that psychedelic drugs open the mind's 'reducing valve'.

Another powerful psychedelic and entheogen found in plants is DMT (N,N-dimethyltryptamine). Sometimes called the 'spirit molecule', DMT induces vivid visual and auditory hallucinations, as well as bodily distortions and out-of-body experiences. American mystic and psychonaut Terence McKenna reportedly said, 'You cannot imagine a stranger drug or a stranger experience', and he had had some very strange experiences.

Smoked in its pure form, DMT acts very fast, the dramatic visual hallucinations and weird sounds coming on almost immediately and lasting only briefly, inviting comparisons with an eight-hour LSD trip compressed into fifteen minutes. For psychologist Ronald Siegel, 'DMT trips are among the most intense drug experiences in the world, and only their brevity makes them bearable' (1992, p. 35). Nick Sand, the underground chemist who first synthesised DMT and discovered that it could be smoked, says, 'What DMT opens up in us is so profound that it is impossible to truly express'. The experience 'has never ceased to amaze me' (Sand, 2014).

Frequent users talk of the DMT 'breakthrough' and take high doses to achieve it, but while some say the breakthrough is a transition into DMT Hyperspace and a very obvious altered state, others say it cannot be described at all.

Swallowing DMT ought to mean a slower and longer-lasting effect, but DMT is quickly destroyed in the stomach by the group of monoamine oxidase (MAO) enzymes that break down adrenaline, dopamine, serotonin, and melatonin. Yet Amazonian shamans have been brewing and drinking DMT in the form of the traditional healing brew called ayahuasca or yagé for hundreds, and possibly thousands, of years. How is this possible?

Ayahuasca is based on the vine *Banisteriopsis caapi* (also called the spirit vine, soul vine, or vine of the dead) mixed with other leaves (e.g. *Psychotria viridis* or *Psychotria carthagenensis*). It is 'one of the most sophisticated and complex drug delivery systems in existence' (Callaway, 1999, p. 256; see also Metzner, 1999). The mixture works because the caapi vine contains MAO inhibitors (the β -Carbolines harmine, harmaline, and tetrahydroharmine), while the other plants contain DMT. If it seems impossible that ancient peoples could have developed this mixture without knowing any chemistry, the truth is probably simpler. When taken alone, the caapi vine has some psychoactive properties – it increases the levels of monoamines such as dopamine and serotonin – so it is possible that this was discovered first and other DMT-containing plants added later.

Traditionally a healing drug, ayahuasca is becoming more popular far away from its original setting, with 'ayahuasca tourism' on the increase. One frequent effect is powerful vomiting, giving the drug another of its common names: the 'vomit drug'. Otherwise, after anything from a few minutes to an hour come a bewildering variety of bodily sensations, transformations, visions, and insights (Metzner, 1999; Shannon, 2002; Luna and White, 2016); we will learn more about the perceptual effects in the section on hallucinations in Chapter 14. A sense of communion with plants and animals is common, and sometimes users feel transformed into the shape and mind of another creature. Contemplation of death is common, as are mystical insights into personal matters and deep existential questions. As with the last drug in this section, drinking ayahuasca is a journey not to be undertaken lightly.

The final drug in this category is often considered to be the ultimate mind-revealing psychedelic: LSD, or d-lysergic acid diethylamide. LSD has a famous

history (Hofmann, 1980; Stevens, 1987). In 1943, Albert Hofmann, a chemist at the Sandoz laboratories in Basel, Switzerland, was working with ergot, a deadly fungus that grows on rye. For eight years he had been synthesising a long series of ergotamine molecules in the hope of finding a useful medicine. Then on Friday 16 April, he synthesised a batch of LSD-25. He began to feel unwell, went home to bed, and experienced a stream of fantastic hallucinations.

Hofmann suspected that, although he hadn't deliberately taken any, the LSD might have caused the hallucinations. Like any chemist wanting to test psychoactive drugs on himself (Shulgin and Shulgin, 1991), he began with what he thought was a tiny dose. On Monday the 19th, at 4.20 in the afternoon, with his assistants present, he took 250 micrograms. At 4.50 he noted no effect, at 5.00 some dizziness, visual disturbance, and a marked desire to laugh. Then he stopped writing, asked for a doctor to be called and, with one of his assistants, set off home on his bicycle.

As he cycled at a good pace he seemed to be getting nowhere. The familiar road looked like a Dali painting and the buildings yawned and rippled. By the time the doctor arrived, Hofmann was hovering near his bedroom ceiling, watching what he thought was his dead body. Instead of the fascinating hallucinations he had had before, this time he was in a nightmare and assumed he would either die or go mad. He did neither, and this first acid trip is now regularly celebrated with re-enactments of his famous bicycle ride (Stevens, 1987). In 2006, the track along which he rode was renamed Albert Hofmann Weg in honour of his 100th birthday. He died in 2008, aged 102.

LSD turned out to be active in tiny doses, and in fact, Hofmann had taken the equivalent of two or three tabs of acid. Like many people since, he had discovered that acid produces the sense of going on a journey, or trip, that can include joy, elation, wondrous hallucinations, deep insights, and spiritual experiences, as well as terrifying horror and despair, and the disintegration of self. To many users, it seems that it opens up the contents of their mind, revealing memories, hopes, fears, and fantasies – both good and bad. This is why there can be bad trips as well as good, and why the term 'psychedelic' is appropriate.

A typical dose of only 100 micrograms induces a trip that begins within half an hour to an hour and lasts anywhere from eight to twelve hours depending on bodyweight, dose, set, and setting. LSD has a chemical structure related to serotonin (5HT) and binds to receptors for serotonin, dopamine, and adrenaline. Although non-addictive, tolerance does build up with frequent use.

The classic work on LSD is *The Varieties of Psychedelic Experience* by Robert Masters and Jean Houston (1967), who observed 206 drug sessions and collected accounts from over two hundred people. They describe bodily distortions, synaesthesia, seeing one's own double, and becoming one with various objects or creatures in the environment, as well as profound religious and spiritual experiences. Early research exploring the use of LSD in therapy yielded extraordinarily positive results (Grof and Halifax, 1977), but research was effectively banned by the drug laws of the 1960s, not beginning again until half a century later. Similarly encouraging results are now again being generated, for example in uses of LSD for the terminally ill (Gasser, Kirchner, and Passie, 2014).

In the first ever placebo-controlled brain-imaging studies, participants had 75 micrograms of LSD intravenously. Functional connectivity increased right across the brain and more local effects coincided with changes in experience. For example, higher cerebral blood flow and greatly expanded functional connectivity in primary visual cortex (V1) correlated strongly with visual hallucinations. Decreased connectivity between the parahippocampus and the retrosplenial cortex correlated strongly with reports of ‘ego dissolution’ and the sense of altered meaning, ‘implying the importance of this particular circuit for the maintenance of “self” or “ego” and its processing of “meaning”’ (Carhart-Harris et al., 2016b, p. 4853). An integrated information-inspired modelling of psychedelic states tries to account for the fact that such states seem to be ‘about more things’ than ordinary waking conscious states, but to involve less systematic organisation and categorisation. From an IIT perspective, this mixture can be explained as a result of increased neural entropy, leading to enhanced cognitive flexibility combined with reduced cause–effect information about all past and future states of the system (Gallimore, 2015).

In other studies, increased connectivity in the temporoparietal junctions correlated with ego dissolution (Tagliazucchi et al., 2016), and decreased functional connectivity within the default mode network (DMN) correlated with less imagery related to the past (Speth et al., 2016). This has implications for treatment of conditions like depression that involve excessive rumination on one’s past, probably mediated by the DMN. Preliminary studies also suggest that LSD, ayahuasca, psilocybin, and other psychedelics can be beneficial for treatment-resistant depression and anxiety without causing harmful side-effects or dependency (Gasser et al., 2014; dos Santos et al., 2016; Frecska, Bokor, and Winkelmann, 2016). The effects include a lasting ‘afterglow’ that can be helpful for people with addictions, and is probably caused by psychedelics’ action on the serotonin system (Winkelmann, 2014).

Psychedelics have changed many people’s lives. Some say they helped solve deep-seated psychological problems, encouraged them to value kindness and love, and inspired them creatively in their work. Many say they were convinced that, for once, they saw things as they really are. But are they right?

Certainly the pioneers of the hippie movement in the 1960s thought so, including Richard Alpert, a young, rich, and highly successful Harvard psychologist who, along with Timothy Leary, Ralph Metzner, and others, had first ‘turned on’ with psilocybin. He then began to find psychology unrewarding and his life empty. Chasing the insights of the drug, he and five others once locked themselves in a building for three weeks and took 400 micrograms of LSD every four hours. It was ‘as if you came into the kingdom of heaven and you saw how it all was and [. . .] then you got cast out again’ (Alpert, 1971, p. 19; Stevens, 1987). He realised how little he knew and went to India to study Eastern religions, and became Baba Ram Dass. Now in his 80s, he is still an active spiritual teacher.

Some people believe that taking LSD has changed them permanently, so it is significant that neural changes in brain entropy after experimental administration of LSD correlated with personality changes two weeks later. ‘Openness to experience’ increased overall, increased more in those who reported ego dissolution during their trip, and changes were still detectable after two weeks. (Lebedev et al., 2016).

Such changes are reminiscent of Huxley's idea of the reducing valve, with users' minds remaining more open for some time. But these effects are hard to classify. Has someone who is changed by taking a major hallucinogen now in a permanently altered SoC, or does their new state become the norm against which other ASCs can be judged?

And are any of these drug-induced ASCs valid, truth-giving, truly spiritual experiences? When people say they transcended duality, did they really see the world in a way that banishes the hard problem and the great chasm? Or are these all just the ramblings of poisoned minds?

In the famous 'Good Friday Experiment', Walter Pahnke, an American minister and physician, gave pills to twenty Boston divinity students before the traditional Good Friday service in 1962: ten received psilocybin and ten an active control (nicotinic acid). Whereas the control group experienced only mild religious feelings, eight out of ten in the psilocybin group reported at least seven of Pahnke's nine categories of mystical experience, developed through work giving LSD to prisoners and the terminally ill: unity, transcendence of time and space, positive mood, sense of sacredness, noetic quality, paradoxicality, ineffability, transiency, and persisting positive changes in attitudes and behaviour (Pahnke, 1963, 1967). Nearly thirty years later, most of the psilocybin group remembered their experiences with clarity and described long-lasting positive effects (Doblin, 1991).

'What I experienced was a God that was inside of me'

(Good Friday participant H.R., in Doblin, 1991, p. 19)

Psilocybin was used again in a double-blind study with thirty-six people who had never had hallucinogens before but who participated regularly in spiritual or religious activities. A high dose of psilocybin or a placebo was given in supportive surroundings over two or three sessions and participants were encouraged to close their eyes and direct their attention inwards. Once again, the drug produced experiences similar to spontaneously occurring mystical experiences, and at a 14-month follow-up, these experiences were considered by volunteers to be among the most personally meaningful and spiritually significant of their lives (Griffiths et al., 2008). (You can find more on Pahnke's work and other exceptional human experiences on our website.)

When we have looked at more ASCs, on the borders of reality and imagination (Chapter 14) and in sleep, dreams, and exceptional experiences (Chapter 15), we shall return to James's question of how to regard them (Chapter 18).

MEDITATION

In Chapter 7, we learned how profound the attentional effects of different kinds of meditation can be. But does meditation induce an ASC? Some definitions imply so: '*Meditation* is a ritualistic procedure intended to change one's state of consciousness by means of maintained, voluntary shifts in attention' (Farthing, 1992, p. 421), and 'meditation can be regarded as a slow, cumulative, long-term procedure for producing an altered state of consciousness' (Wallace and Fisher, 1991, p. 153).

Some forms of meditation, such as transcendental meditation (TM), do emphasise the importance of achieving altered states, but others do not. In Zen, the aim of

practice is not to achieve an ASC or reach any other goal. Rather, meditation itself becomes the task (Watts, 1957). 'Enlightenment and practice are one', claimed the thirteenth-century Zen teacher Eihei Dogen.

Even so, some Zen Buddhist practitioners may have dramatic, if temporary, *kenshō* (awakening) experiences, including glimpses of the true nature of mind, experiences of emptiness, and great flashes of understanding leading ultimately to the 'dropping-off of body and mind' (Chapter 18). As the story goes, Dogen was sitting one morning in meditation when his master reprimanded a dozing monk, urging him to wake up and work harder, saying 'To realize perfect enlightenment you must let fall body and mind'. In that moment, Dogen achieved full awakening or liberation (Kapleau, 1980). Something had clearly changed, and changed dramatically, yet it is said that enlightenment itself is not a state of consciousness.

So does meditation induce ASCs? According to Tart's subjective definition, it does, because people feel that their mental functioning has been radically altered. One sceptical theory is that meditation is nothing more than sleep or dozing. Neuropsychiatrist Peter Fenwick (1987) (see the website for a brief biography) showed that EEG profiles in meditation are not the same as sleep or drowsiness, yet many meditators take microsleeps during meditation, and in one study TMers slept as much as a third of the time while meditating (Austin, 1998). Since naps are known to reduce anxiety and depression and improve cognitive ability, this might explain some of the claimed effects of meditation. Yet meditators themselves say they can easily distinguish between deep meditation and sleep. One interpretation is that meditators learn, with inevitable slip-ups, to hold themselves at that interesting transitional level between sleep and wakefulness (Chapter 15).

Perhaps the most extraordinary claim for ASCs achieved through meditation appears in the early Buddhist teachings of the Theravadan tradition. The jhanas are a series of eight increasingly absorbed states said to be reached through deep concentration applied in a series of graded steps (Chapter 18). The first jhana involves raising a kind of 'energy' that the ancient suttas refer to as *pīti*, and may be equivalent to the kundalini energy described in some yoga and other traditions. This can come in a rush of shaking, vibrating, noises, and hot flushes and is maintained by attending to the sense of glee or joyfulness that fills the whole body. The skill is then to drop down from this hyper-excited state into a happy but calmer state, and then to equanimity, converting the *pīti* into a gentler 'energy' called *sukha*, and finally into a deep, emotionally neutral state without thoughts.

These first four states are firmly rooted in the body, but the next four are referred to as the 'immaterial' or 'formless' states. Limitless, infinite, or boundless space entails continuous expansion, until only endless space remains as the object of attention. Cannabis can have this same effect of expanding into space with no bodily sense remaining (Tart, 1971; Blackmore, 2017). Beyond this lies the sixth jhana of infinite consciousness and beyond that further indescribable states.

From these descriptions, the states are clearly meant to unfold in a specific order, and meditation teacher Leigh Brasington speculates that the techniques

'a d[iscrete]-ASC in meditation would be expressed in a discrete state of brain networks'

(Thomas and Cohen, 2014, p. 5)

amount to controlled self-stimulation of the reward system. This begins with a flood of dopamine, leading to increased noradrenaline and then to endorphins, each neurotransmitter accounting for the various emotions and sensations of the first three jhanas. Finally, the opioids fade, leaving the neutral state of the fourth jhana (Brasington, 2015). Although speculative, these ideas can be and have been tested. When Brasington meditated inside scanners, the shifts from each state to the next could be seen with both EEG and fMRI (Hagerty et al., 2013). In further studies, increased activation of the nucleus accumbens was found to correspond to the extreme joy of the first jhana, which makes sense because this is part of the dopamine/opioid reward system. If it turns out that the jhanas are a naturally occurring sequence of brain-based states, they might provide an excellent example of Tart's 'discrete states of consciousness' (d-ASC).

A methodological review of meditation research (Thomas and Cohen, 2014) takes the term d-ASC to imply that we should look not just for correlations between the phenomenology and the physiology (i.e. the experience and the neural activity), but for 'recognizable isomorphism' between them: 'Thus, a d-ASC in meditation would be expressed in a discrete state of brain networks, observable as a change in the dominant network of functional connectivity between brain regions, from a defined baseline state' (p. 5).

'Zen training means brain training'

(Austin, 1998, p. 11)

Note that although this suggestion might seem similar to the search for the NCCs, there is an important difference: here the idea is to correlate specific d-ASCs with specific changes in brain activity, not to search for the correlates of conscious versus unconscious states or for 'consciousness itself', which may or may not exist. The hypothesis is also stronger than many, with a higher threshold for what counts as 'alteration', and a requirement that we define the state from which the alteration departs. The authors also urge a multidimensional approach to studying ASCs in meditation, studying the person (characteristics of the meditator), practice (the specific meditative style), place (the experimental situation and wider geographical and cultural context), phenomenology (the meditator's experiences), and psychophysiology (including documentation of methods). We return to these considerations in Chapter 17.

One possibility is that states reached by novice meditators may overlap with states occurring outside meditation practice (e.g. relaxation), even if they occur more reliably and last longer, but that advanced meditators may reach states that are unique to meditative practice (Fell et al., 2010), perhaps because meditating gradually changes the neural structures of the brain. This is suggested, for example, by the combination of increased synchronicity in both low-frequency oscillations and gamma activity in experienced meditators – gamma activity normally being *reduced* in relaxation and sleep.

What it really means to claim 'uniqueness' for any experience – let alone a complex set of experiences found during as wide a set of activities as meditative practice – remains unclear, although the jhanas may provide more clearly 'unique' states than are found with more general meditation practices. Maybe it makes sense to hedge our bets slightly and talk about 'meditation-related states of consciousness' (Fell et al., 2010) rather than ASCs, to avoid commitment to a strong view of what is altered relative to what.

MENTAL ILLNESS

The term ASC seems to be vague enough that it can be attributed to almost any identifiable shift in experience, including fluctuations in ordinary wakefulness like daydreaming, drowsiness, hypnagogic states, dreaming, and sleep; these would all be grouped under the heading of 'spontaneously occurring ASCs' (Vaitl et al., 2005). Then there are ASCs induced by extreme environmental conditions like heat and cold, altitude and microgravity, as well as those induced by starvation or orgasm; these might fall into the categories of physiologically induced ASCs. Some kinds of illness are capable of inducing ASCs, including those that cause sleep deprivation or oxygen deprivation, fever, or seizures, while psychologically induced ASCs might range from rhythmic trance to sensory deprivation to bereavement – and perhaps hypnosis (see Concept 13.2).

A final case we will consider now is that of mental ill health. This brings some of the problems with the concept of ASC into sharper focus, as well as raising questions about personal identity and responsibility, topics sometimes covered in the field of philosophical psychopathology (Gennaro, 2017).

The first point to make about mental illness is that it is never solely mental. All psychological disorders involve feedback loops between thought patterns, emotions and moods, behaviours, and bodily states. This already makes a neat classification into 'psychologically' and 'physiologically' induced ASCs impossible: where do we classify anorexia, or self-harm, or anxiety disorders involving repetitive obsessive-compulsive activities?

The second thing to note is that one of the factors which helps sustain mental ill health is the difference between the nature of experience while ill and when healthy; this can make it hard to remember, imagine, or believe in the reality of a state of consciousness other than the pathological one, which can reduce the motivation to seek help or persist in recovery. In this sense, mental illness seems an obvious candidate for inducing ASCs. *What it is like to be*



CONCEPT 13.2

IS HYPNOSIS AN ASC?

The term 'hypnosis' comes from the Greek *hypnos*, for sleep. Nineteenth-century researchers believed that under hypnosis people fell into a state of somnambulism, or sleep-walking. Others argued against a special hypnotic state and attributed all the effects to suggestion and imagination. This argument turned into the twentieth-century battle between 'state theorists' and 'non-state theorists'.

Applying Tart's definition, we should easily accept that hypnosis is an ASC, because hypnotic subjects often feel that their mental functioning is radically different from normal. However, the traditional view of the hypnotic state carries far more contentious implications, implying a dissociation between different parts of the mind. This idea reappeared in the 1970s with the neo-dissociationist theory of American psychologist Ernest Hilgard.

Hilgard (1977) argued that an executive ego normally directs multiple control systems, but in hypnosis the hypnotist takes over this function, making the subject feel that his or her actions are involuntary and suggested hallucinations are real. Against this, non-state theorists argue that hypnotic subjects are playing a social role to please the experimenter, using imaginative strategies to comply with suggestions, or simply faking it (Spanos, 1991; Wagstaff, 1994).

In support of dissociation, Hilgard discovered the phenomenon of the hidden observer. When a hypnotised subject with her hand in freezing water claimed not to feel pain, Hilgard suggested he could talk to 'a hidden part of you', whereupon the person (i.e. the hidden part) described the pain she was feeling. Critics responded by demonstrating that the hidden observer could be made, by appropriate suggestions, to feel less pain instead, implying that it was all just suggestion after all (Spanos, 1991).

Crucial experiments compared 'real' hypnotised subjects with controls asked to fake or imagine being hypnotised, arguing that if controls show the same phenomena as 'really hypnotised' subjects, the idea of a special hypnotic state is redundant. Many experiments have shown no differences, but there are interesting anomalies. Using 'trance logic',

hypnotised subjects seem able to accept illogicalities in a way that fakers cannot. For example, when asked to hallucinate a person who is actually present, they may see two people while simulators tend to see only one. When regressed to childhood, they may describe feeling simultaneously grown-up and child-like, while simulators claim to feel only like a child. Similar trance logic can be seen in some drug states, dreams (Chapter 15), and mystical experiences (Chapter 18).

In the mid-1990s, British psychologist Graham Wagstaff concluded that 'in over one hundred years we seem to be no further forward in deciding whether there is an altered state of consciousness that we can call "hypnosis"' (1994, p. 1003) and the debates continued (Kallio and Revonsuo, 2003, with peer commentaries and response, 2005). Research has found changes in neural activity, such as changes in anterior cingulate related to reduced pain in hypnosis (Faymonville et al., 2000), but the relationship between these changes and the hypnotic induction remains unclear (Mazzoni et al., 2013).

A veteran of hypnosis research, Irving Kirsch, has called the disagreements 'needlessly vociferous' and lacking in 'heuristic value', arguing that the more interesting and important question is 'how highly suggestible subjects are able to experience sometimes profound subjective alterations' (2011, p. 353). He concludes that what produces heightened suggestibility may be the person's perception of being in an altered state, rather than some state *in itself*. This way of thinking challenges the idea that a 'state' of consciousness can be distinguished from what the person experiencing it wants or believes it to be.

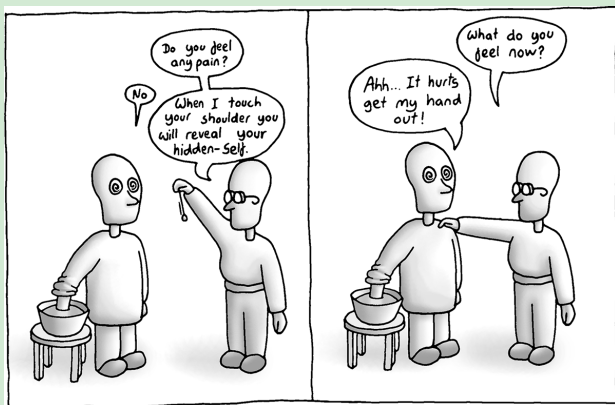


FIGURE 13.8 • The hidden observer. Although a hypnotic subject with his hand in freezing water may claim to feel no pain, Hilgard discovered that by giving an appropriate signal, he could talk to a hidden part of the person who was in pain. This forms part of the evidence for neo-dissociationist theories.

me is profoundly altered in illness versus health, to the extent that I might even stop believing I could exist without the illness.

[T]o me, 'Emily' became nothing more or less than anorexic Emily. My blank or distraught or irritable or fragile moods, my need for routine and privacy, my slight figure, my lack of friends and my worship of academic achievement, all seemed like innate parts of me, and there seemed no reason to believe that eating breakfast or lunch would make a difference to any of them. The extent to which 'I' was the product of years of malnutrition and the rigid, ritualised mental life and physical limitations that malnutrition itself created was not something I was capable of comprehending, since to do so I would have had to imagine my life as otherwise than it was – and I had neither the ability nor the desire to do that. It was a perfect vicious circle: the anorexia had become so completely what I was that I couldn't see how completely it had taken over 'Emily', nor could I therefore have any motivation to try to find her again.

(Troscianko, 2012, p. 242)

But does this mean that we should think of the illness as itself an ASC, or as something which brings about an ASC (or multiple ASCs)? For Antti Revonsuo and colleagues (2009),

the definition of an ASC refers to the temporary (as opposed to permanent) nature of alterations in the representational mechanisms of consciousness. The altered state commences at some specifiable time-window, and the normal state of consciousness and brain returns at some later time.

(p. 196)

This means that if psychosis such as schizophrenia were a permanent pathological state, it could not be an ASC, but temporary episodes within it could be.

Any neat distinction between permanent and temporary, illness and episode, is easy to question: does it really make sense to separate the ongoing distortions brought about by chronic semi-starvation in anorexia, for example, from the

shorter-term effects of acute fasting? There are differences, to be sure, but why does one count as an 'altered state' and the other not? In other kinds of illness, like bipolar disorder, discrete episodes of psychosis may come and go, but for illnesses in which transitions between moods and other cognitive states are more continuous (like some forms of depression, say), would there be no 'altered' consciousness because all the boundaries are blurred and the timescales protracted?

As their definition shows, Revonsuo and colleagues' answer is that the alteration in an ASC is an alteration not to consciousness per se, but to the representational relationships between consciousness and the world, with the 'neurocognitive background mechanisms of consciousness' producing 'misrepresentations such as hallucinations, delusions, and memory distortions' (2009, p. 187). This argument is intended to make 'normal' and 'altered' objectively definable in terms of the accuracy of information being transferred from 'world' to 'consciousness'. But given all the arguments about mental and neural representation we explored in Chapter 3, it is unclear whether accuracy of our information about the world can reliably help separate 'unaltered states' from ASCs – or whether it even makes sense to say that there is a representational relationship between consciousness and the world.

As well as relying on such concepts as 'conscious representation in the brain' and 'content in phenomenal consciousness', this line of thinking ultimately reduces an ASC to its neural correlates:

to objectively determine the presence of an ASC, one must show that the background mechanisms of conscious representation in the brain are altered in a way likely to lead to (globally and temporarily) misrepresentational content in phenomenal consciousness.

(2009, p. 196)

Clearly Revonsuo and colleagues are urging an objective definition of ASCs, with all the problems that entails. Changes in neural activity are easily identified in both illness and recovery, with one study on psychosis, for instance, finding that changes in connectivity in brain different areas predicted improvement in psychotic or emotional symptoms following cognitive behavioural therapy (Mason et al., 2017). But this does not necessarily mean we can or should define psychosis in terms of those connective patterns.

The case of mental illness also raises again that nagging question about the baseline from which 'alteration' is made. If any illness can be accompanied by or induce an ASC, health is presumably the baseline. But how do we define that? For the person concerned, the differences between mental illness and health are tangible and life-defining. And for any mental illness there are concrete ways of operationalising the kinds of suffering it involves, for the unwell person and sometimes also for other people. The difficulties come when we try to pin down precise points of transition, in time or in quality of life: where does dieting stop and an eating disorder begin, for instance, or exhaustion shift into chronic fatigue?

In general terms, we could define health simply as the absence of illness, but if we try to do better than that, we may find ourselves gravitating towards ideas

'it is not the state itself that is producing heightened suggestibility but rather the person's perception of being in an altered state'

(Kirsch, 2011, p. 359)

'psychiatric diagnostic labels [. . .] should not be classified as ASCs or not. Only the psychotic episodes, had by any kind of patients, can be ASCs.'

(Revonsuo et al., 2009, p. 201)

'an ASC should not be defined as an altered phenomenal state of consciousness, but an altered representational state'

(Revonsuo et al., 2009, p. 196)

like wellbeing, calm, compassion, openness – yet these are culturally variable and hard to defend or define in detail. A representational view like Revonsuo and colleagues’ does not help us here, because perfect representational matching between consciousness and world may sound like sanity, but it doesn’t sound much like health. And in any case, we know (Chapter 3) that representations of the world are optimised not for accuracy of representation but for efficiency. So can I ever say with confidence that my experience of health would not count as ‘altered’ compared to yours? And if I recover from long mental illness, the health I return to will of course not be the same as the health I experienced before the illness.

That was the strange thing, that one did not know where one was going, or what one wanted, and followed blindly, suffering so much in secret, always unprepared and amazed and knowing nothing; but one thing led to another and by degrees something had formed itself out of nothing, and so one reached at last this calm, this quiet, this certainty, and it was this process that people called living.

(Virginia Woolf, *The Voyage Out*, 1915)

An interesting twist here is that mindfulness meditation and several kinds of psychoactive drug seem to be effective in treating mental illness. That is, techniques often used for inducing altered states can also be used to cancel others out. ‘Mindfulness-based cognitive therapy’ is helpful for conditions including depression (Piet and Hougaard, 2011). Meanwhile, ‘microdosing’ of psychedelics such as LSD is increasingly used as self-treatment for various mood disorders (Maughan, 2017), and we summarised earlier in the chapter some of the research evidence for the therapeutic use of LSD as well as psilocybin and MDMA. In some cases, the drug being used may function simply to bring some aspect of brain function back to normal, for example by stimulating underactive serotonin receptors, as LSD does. But the chemical and neural structures involved are complex, as are individual histories and environments. So it seems odd to say that an ASC (such as PTSD, or some episodes in the experience of PTSD) could simply be negated by a substance (like MDMA) or practice (like mindfulness meditation) that is usually considered to induce a different ASC. This kind of mathematical cancelling-out (alteration¹ + alteration² = baseline) seems implausible.

There are so many ways of establishing and then questioning an ‘alteration to consciousness’ that the very concept starts looking rather like ordinary life: our experiences are never exactly the same even two minutes running, and once we start trying to qualify or quantify what counts as a proper alteration, we quickly find ourselves on shaky ground. There is a strong argument to be made that ‘normality’ still needs as much investigation as the ‘alterations’ from it, and while much of the research covered in this book could be seen as attempting this, we may also need to develop more contextually sensitive methods. One option would be an ethnological approach to mapping the continuities and variations in what different people from different cultures take to be normal.

ASCs have always had negative connotations (being associated with the different, the strange, the abnormal, the irrational, the pathological) in tension with their positive ones (the wonderful, the insightful, the life-altering). Some ASCs clearly should be called pathological: if they seriously impair quality of life for the person experiencing them or for those around them, for example. But the literature on drug use, sprinkled with phrases like 'relapse to cannabis use' (Crean, Crane and Mason, 2011), makes clear that normative judgements are in play far more widely, imposing pathology where there may be none. Some unusual experiences of other kinds may also be inappropriately pathologised as mental illness instead of being understood as instances of personal transformation or post-traumatic growth or simply human variation (out-of-body experiences, covered in Chapter 15, are just one example of this). On the other hand, destigmatising mental illness is a process still very much ongoing, and that process requires a willingness to draw boundaries between health and illness even when there is no clear organic cause.

STATES OF CONSCIOUSNESS

We have explored at length the question of whether talking about *altered* states of consciousness makes sense. To conclude, it is worth asking whether *state* is the most helpful word to use. It seems obvious what is meant by a state of consciousness, but we should bear in mind that to speak of a state is to assume there must be something which is in that state (or condition). And what is that something? If instead of a thing called consciousness (e.g. a container with contents) we imagine a process of attribution after the fact (as in the multiple drafts model), then there is nothing to be in a state or not in a state.

In the long debate about whether hypnosis does or does not induce an ASC, one attempt at a resolution was simply to reduce the term *state* to a mere label, 'a kind of shorthand, with no causal properties or defining features associated with it' (Kihlstrom, 1985, p. 405). Making the central concept so meaningless would effectively end the whole debate. Psychologist Irving Kirsch (1997, p. 98) noted that in response, various euphemisms for state, like 'special process', arose to disguise the fact that the state debate was still happening, even though multiple, often closely related positions had emerged, and the whole thing was no longer getting anyone very far.

You could say all this is just semantics. Kirsch argues not:

if hypnosis is a state, like sleep or intoxication, then establishing its essential characteristics is an important task for hypnosis researchers. Conversely, if the state hypothesis is false, these questions are meaningless and research should be directed elsewhere.

(1997, p. 97)

But we can argue against states of consciousness without abandoning the study of consciousness. In the case of hypnosis, we can still study the beliefs, expectations, imaginative strategies, and everything else that makes hypnotised experiences unusual.




FIGURE 13.9 • Kirsch (2011) suggests that theoretical positions are on a continuum rather than a dichotomy. Might this idea be helpfully applied to the contentious question of whether hypnosis is an altered state or not?

of conscious experience might contribute to the mystery of 'consciousness itself'. It would also mean that we ignore at our peril forms of experience we may fear, disapprove of, or have no interest in, since any defence of a single baseline 'normality' looks decidedly shaky.

Maybe both state and alteration are red herrings. And this would mean there is even more work to do than we thought in making sense of what all the varieties

READING



Doblin, R. (1991). Pahnke's 'Good Friday Experiment': A long-term follow-up and methodological critique. *The Journal of Transpersonal Psychology*, 23, 1–28.

Provides lots of methodological detail on Pahnke's original experiment on psilocybin and spirituality, including reports from participants and ethical considerations.

Jay, M. (Ed.) (1999). *Artificial paradises: A drugs reader*. London: Penguin.

Contains brief extracts from de Quincey, Huxley, Freud, Davy, Hoffman, Shulgin, James, Siegel, Leary, Tart, Grof, and many others.

Kallio, S., and Revonsuo, A. (2003). Hypnotic phenomena and altered states of consciousness: A multi-level framework of description and explanation. *Contemporary Hypnosis*, 20(3), 111–164. Peer commentary and authors' response in Gruzelier, J. (2005). Altered states of consciousness and hypnosis in the twenty-first century. *Contemporary Hypnosis*, 22(1), 1–54.


An account of ASCs aiming to address the state/non-state question about hypnosis, with a programme for future research embracing multiple levels of description.

Tart, C. T. (1972). States of consciousness and state-specific sciences. *Science*, 176, 1203–1210. See also Tart's 2015 blog post at <http://blog.paradigm-sys.com/state-specific-sciences-altered-state-origin-of-the-proposal/>

The article proposes a new science conducted by people in specific SoCs, including the potential and the pitfalls. The blog post describes how the idea came about and how others responded.

Thomas, J. W., and Cohen, M. (2014). A methodological review of meditation research. *Frontiers in Psychiatry*, 5, article 74.

Argues that a broader scope is needed if meditation research is to satisfactorily answer questions like whether meditation induces ASCs.



Reality and imagination

CHAPTER

FOURTEEN

‘When they saw where she lay, which they had not done till then, they showed no objection, and stood watching her, as still as the pillars around. He went to the stone and bent over her, holding one poor little hand; her breathing now was quick and small, like that of a lesser creature than a woman. All waited in the growing light, their faces and hands as if they were silvered, the remainder of their figures dark, the stones glistening green-gray, the Plain still a mass of shade. Soon the light was strong, and a ray shone upon her unconscious form, peering under her eyelids and waking her.

“What is it, Angel?” she said, starting up. “Have they come for me?”

“Yes, dearest”, he said. “They have come.”

“It is as it should be”, she murmured. “Angel, I am almost glad – yes, glad! This happiness could not have lasted. It was too much. I have had enough; and now I shall not live for you to despise me!”

She stood up, shook herself, and went forward, neither of the men having moved.

“I am ready”, she said quietly.

[. . .]

Upon the cornice of the tower a tall staff was fixed. Their eyes were riveted on it. A few minutes after the hour had struck something moved slowly up the staff, and extended itself upon the breeze. It was a black flag.

“Justice” was done, and the President of the Immortals, in Aeschylean phrase, had ended his sport with Tess. And the d’Urberville knights and dames slept on in their tombs unknowing. The two speechless gazers bent themselves down to the earth, as if in prayer, and remained thus a long time, absolutely motionless: the flag continued to wave silently. As soon as they had strength, they arose, joined hands again, and went on.’

(Thomas Hardy, *Tess of the d’Urbervilles: A pure woman faithfully presented*, 1891)

What do you feel when you read this passage? How do you imagine the two scenes, with Tess and after her death? Do you notice any changes in your body or your awareness of your surroundings as you imagine? Do they come in response to particular phrases or sentences? Maybe you have read the whole novel, and many other memories of Tess and Angel come to mind. Maybe you have experienced shame or bereavement in ways that heighten your emotion now. Or maybe the passage seems too sentimental and this description leaves you cold. But for us at least, the imaginative experience has a visceral reality. We know none of these characters exist; and their situation may have little outwardly to do with our own lives. And yet, what we imagine is real in the sense that it has effects on us: reading creates an experience, and the *experience* is real. Thinking this way can make us deeply confused about the difference between reality and imagination.

Perhaps we should be confused. Let’s take another example. Suppose you walk into your kitchen and see your black cat on the chair. You look again and realise that it’s actually a friend’s pullover, left in a heap, with one arm dangling. The strange thing is that if you had not looked again, you could have described how the cat was sitting, which way its ears were pointing, and how its tail hung down off the seat. You may say that the pullover was real and the cat was imagined, but now consider the same thing happening when the cat is actually there. In a brief glimpse you could not have taken in all those details, and yet they were mostly correct. You noticed no gap corresponding to the blind spot on your retina (Chapter 3). You saw a whole cat even though its hind legs were half hidden by the back of the chair. Was what you saw real or imagined?

Other examples that confront us with this question include apparently simple perceptual events like ‘watching the sun go down’ at ‘sunset’. You know that the earth is moving relative to the sun, not vice versa, and yet your experience is of the sun sinking down beneath the horizon. Or take a well-known case of a social illusion: ‘the dress’. The most popular internet meme of 2015, generating 840,000 views per minute and 4.4 million tweets in 24 hours, was a photograph originally posted on Tumblr of a dress that to some people looks like blue and black stripes and to others white and gold. The actual colours were eventually confirmed to be black and blue, though the manufacturer produced a one-off white-and-gold version for a charity auction. Vision scientists became intrigued by this striking example of a bistable colour stimulus with a very low switch rate. The *Journal of Vision* has an ongoing special topic (Allred et al., 2017) devoted to exploring the phenomenon: a full-scale twin study was used to determine the contributions of genetic and environmental factors, and other articles investigated the effects of sensitivity to contextual cues about the illumination of the scene, and the relative contributions of stable traits of the visual system versus one-shot learning effects.

Some coverage of #thedress suggested that it might prompt an ‘existential crisis’ about the nature of vision and reality, and many people described having arguments about it, feeling freaked out by it, wondering ‘is my brain tricking me?’ and ‘is it a conspiracy?’

These phenomena bring us back to all the familiar philosophical problems involved in what it means to see, and to the central problem of consciousness: the difference between the objective and subjective worlds. In particular, they return us to the idea that we may be so seriously mistaken about consciousness that we should think of it as a grand illusion or as something that does not exist in the form we usually take it to (Chapter 3). Perhaps it may help to learn about other strange experiences that hover between reality and imagination.

REALITY DISCRIMINATION

In everyday life, we discriminate ‘real’ from ‘imagined’ all the time without noticing the skill involved. That is, we distinguish our own thoughts from what we assume to be a public reality independent of those thoughts – a skill called reality monitoring (Johnson and Raye, 1981) or reality discrimination. Experiments in which people are asked to see or hear some stimuli, and to imagine others, show that many different features can be used for the purpose of discrimination, including how stable, detailed, or vivid the experiences are, and whether they can be voluntarily controlled. One study (Garrison et al., 2017) presented participants with either complete or incomplete well-known word pairs (‘Laurel and Hardy’ or ‘Eggs and . . .’) and tested how well they remembered which words were actually presented and which needed completing imaginatively: visual presentation resulted in better reality monitoring than auditory presentation, and speaking the words out loud worked better than internally verbalising (‘thinking’ about) them.

By and large, mental images are less richly detailed, less stable, and more easily manipulated than perceptions. So we don’t usually confuse the two. We can, however, be tricked. In her century-old classic experiment, Cheves Perky (1910) asked participants to look at a blank screen and to imagine an object on it, such as a tomato. Unbeknown to them, she back-projected a picture of a tomato onto the screen, and gradually increased the brightness. Even when the picture was bright enough to be easily seen, the participants still believed that they were imagining it. This effect is the reverse of hallucinations, in which we think something is there when it isn’t. In this case, Perky’s participants were tricked into thinking there was nothing there when there was. Similar effects have been found since, and show that reality discrimination is affected by whether we expect something to be real or imagined.

Distinguishing memories of events that really happened from events we have only imagined is particularly difficult, and its failure results in false memories – that is, convincing ‘memories’ of events that never actually happened. These can be created when we tell the same story many times, with slight variations, and then remember the last version we told. The latest version retroactively interferes with the original memory. False memories can also be created when a family story keeps being told or a photograph from childhood convinces you that you can remember that day on the beach. And they can have lasting effects on behaviour. For example, when people were told that they had had positive or negative

experiences with particular foods in their childhood, their expressed preferences for those foods, and their eating of them, was affected months later (Geraerts et al., 2008).

False memories can also be deliberately created by asking leading questions that encourage someone to invent an answer concerning something they never experienced. In a famous example, psychologist Elizabeth Loftus showed participants a film of a traffic accident and asked how fast the cars were going when they either smashed into, collided with, bumped, or hit each other. Those who heard 'smashed' gave higher speed estimates and a week later were more likely to 'remember' broken glass in the film when there was none (Loftus and Palmer, 1974).

False memories are most problematic when people 'remember' sexual abuse that never happened, or identify suspects they never saw (Loftus and Ketcham, 1994). There have been tragic cases in which therapists allegedly recovered repressed memories of sexual abuse under hypnosis and convinced their patients that the events really happened when they did not. The processes involved are still poorly understood and there is no agreement over whether memory repression or suppression, or 'motivated forgetting', really happens or not (McNally, 2012; Staniloiu and Markowitsch, 2014). In such cases, it is often possible to find out the truth of what happened by some kind of independent verification, but this still does not mean that there is a sharp dividing line between 'real' memories and 'false' memories. So how come we can normally be so confident about our own memories?

Real memories tend to be more detailed and more easily brought to mind than false memories. Sometimes real memories can be identified because we can put them in context with other events or remember when and how they happened – a skill called source monitoring. This is not important for learning skills and facts. For example, you may reliably and correctly remember the speed of light, the capital of Germany, and the name of the man next door, without needing to remember when or where you learned them, but for autobiographical memory the context is important. If the memory of an event in your life is detailed and plausible, and fits with other events in time and place, then you are more likely to judge that it really happened.

We probably all hold false memories, and even valid memories may consist of accurate elements mixed with plausible concoctions and embellished with invented details, because autobiographical memory is nothing like a static archival device by which memories can be encoded, stored, and retrieved; rather, remembering is a process of active reconstruction shaped as much by our current goals and priorities as by the realities of the past (Conway, 2005). Striking overlap has been found in the neural activity associated with remembering past experiences and imagining possible future experiences, and parallels have also been observed in how memories and future imaginings are structured and experienced in relation to factors like emotion, level of detail, and psychopathology (Schacter et al., 2012). Memory also has profoundly social aspects in the sense that past experiences are not easily divided into 'mine' and 'yours', and every social act of recollection changes the next (Saunders, 2014). Some have even suggested that a memory about something in your past is not a belief about the past, but constructed in order to justify, to yourself and others, *why you hold* a particular

'human memory can be remarkably fragile and even inventive'

(Geraerts et al., 2008, p. 749)

belief – something we have to do in many social negotiations about entitlements and obligations (Mahr and Csibra, 2017).

Even the concepts we often use to think about memory can lead us astray: we tend to use the word ‘vivid’ to describe powerful memories, but when we talk about vividness, are we referring to the level of detail or accuracy of the memory or its intensity (Jajdelska et al., 2011)? There is evidence that a strongly emotional sense of being ‘brought back’ to a past time can make people assume there is more detail in their memories than is actually the case (Herz and Schooler, 2002). Language plays many other roles in the fluidities of human memory. False memories can be created by presenting people with short narratives of plausible but false childhood events mixed in with a few real ones; during interviews, people need little encouragement to search for relevant thoughts, images, and feelings and so create the ‘memory’. Doctored photographs can have the same effect: in a study which asked participants to describe all they could remember about the family events pictured, half of the twenty participants had talked themselves into believing the memory was real, with their emotional involvement increasing across the three interviews (Wade et al., 2002).

The dividing line between real and unreal is obviously blurred in the context of memory. The division is particularly interesting when it concerns experiences for which there can be no public corroboration, including dreams, fantasies, and hallucinations. Did I *really* feel moved by a woman’s death? What does this question mean?

HALLUCINATIONS

DEFINING HALLUCINATIONS

The term ‘hallucination’ is not easy to define, although some rough distinctions can be helpful if not applied too rigidly. Hallucinations were distinguished from illusions early in the nineteenth century, on the basis that hallucinations are entirely ‘internal’ whereas illusions are misperceptions of ‘external’ things. Illusions include familiar visual illusions such as the Müller-Lyer, Ponzo, or Café Wall illusions (see Chapter 3), as well as misperceptions like seeing a pullover as a cat. By contrast, hallucinations are perceptual experiences not elicited by an external stimulus. This distinction is still used (e.g. Waters et al., 2014), but there is no clean dividing line. For example, imagine that someone sees the ghost of a headless monk float across the altar in church. We might say that there was nothing there and that the monk was a hallucination, or alternatively that a faint swirl of candle smoke or incense was misperceived and that the monk was an illusion.

True hallucinations are sometimes distinguished from pseudo-hallucinations, in which the person knows that what is seen or heard is not real. For example, if you heard a voice telling you that the thought-police were coming to get you, and you believed they were, you would be suffering a true hallucination, but if you heard the same voice as you were nodding off over your computer, and realised you were working too late, that would be a pseudo-hallucination. One problem with this distinction is that if taken too literally there must be very few true hallucinations. Even with a double dose of LSD, most people still know that the arms of the enormous monster threatening to engulf them are really the branches of a tree,

and when a uniformed naval officer appears at the wheel in the midst of a storm hundreds of miles from shore, the exhausted lone sailor knows that no one else can really be aboard the ship.

A final distinction is between hallucinations and mental imagery. Hallucinations are sometimes distinguished from imagery by their resemblance to publicly shared perceptions rather than private thoughts, or by their uncontrollability. If we voluntarily imagine a tropical beach with the sound of waves lapping on the sand, this is usually called imagery, but if the vision forces itself on our mind and won't go away it would be called a hallucination. But even this distinction is unclear. For example, the images that come on the borders of sleep (Chapter 15) are usually called 'hypnagogic imagery' rather than 'hypnagogic hallucinations', although they are not voluntary or easily controllable. Imaginative sensory-like experiences that occur while we read fiction are normally called imagery too, even though they are guided by the text and your imaginings may be intrusive and hard to set aside, even after you put the book down. So, rather than try too hard to delimit these categories, some prefer to think of a continuum with true hallucinations at one end and imagery at the other. But even this may not help if there are multiple dimensions involved.

These distinctions are discussed by British psychologists Peter Slade and Richard Bentall (1988), who propose this working definition of a hallucination:

Any percept-like experience which (a) occurs in the absence of an appropriate stimulus, (b) has the full force or impact of the corresponding actual (real) perception, and (c) is not amenable to direct and voluntary control by the experienter.

(p. 23)

This too has its problems, especially with point (b). What does the 'full force or impact' mean when applied, for example, to a ghostly human figure seen climbing dimly lit stairs? Many such hallucinations are described as fleeting and the figures as transparent, but transparent people do not exist, so there is no obvious 'actual (real) perception' to compare this experience with.

PREVALENCE OF HALLUCINATIONS

One of the first attempts to study hallucinations in the general population was the Society for Psychical Research's 'Census of hallucinations' in the late 1800s (Gurney, Myers, and Podmore, 1886; Sidgwick et al., 1894). This was a time when 'spirit mediums' were holding séances right across Europe and the USA, some in complete darkness, with spirit voices emanating from luminous floating trumpets, music mysteriously playing, and touches and cold breezes being felt. Sometimes a translucent, greyish substance called ectoplasm was exuded from the bodies of certain mediums, and even 'materialised' into the bodily form of spirits (Gauld, 1968). Many were caught cheating, and mediums who wanted to enhance their act could easily purchase muslin drapes, trumpets, luminous paint,




FIGURE 14.1 • On LSD trips the floor can turn into a carpet of snakes, cars into space ships, and trees into monsters. But typically the tripper still knows that the monster is really a tree, and is therefore technically having a pseudo-hallucination, not a true hallucination.



FIGURE 14.2 • In the heyday of spiritualism mediums were tied up inside a ‘cabinet’ while the ladies and gentlemen watched. In a deep trance, they claimed to exude ectoplasm from various orifices of the body and so create fully formed spirits that could move around the room, touching the astounded sitters and even answering their questions.

‘Among those implications none can be more momentous than the light thrown by this discovery [of telepathy] upon man’s intimate nature and possible survival of death.’

(Myers, 1903, i, p. 8)

and special chairs from which they could easily escape, even if sceptical observers tied them to the chairs with ropes.

Even without any cheating, the traditional darkened séance room provides ideal conditions for complex interactions between imagination and reality involving illusions, hallucinations, motivated errors, and criterion shifts. British sceptic and parapsychologist Richard Wiseman recreated similar conditions in fake séances in which an actor suggested that a table was levitating when it was not. Small objects were painted with luminous paint and a hidden assistant moved them about in the darkness using a long stick. About a third of participants reported afterwards that the table had moved, including more believers than disbelievers (Wiseman et al., 2003).

Spiritualism was ignored by most scientists but appealed to people who felt threatened by the materialism of Victorian physics and the radical new ideas of Darwinism that seemed to undermine the special status of humanity. After all, if spirits of the dead could appear and speak, then materialism must be false. (See the website for more on the paranormal in the 19th century.)

It was in this context that, in London in 1882, the Society for Psychical Research (SPR) was founded by a small group of highly respected scientists and scholars to examine these and other claims of psychic phenomena, and one of their first achievements was the Census of Hallucinations. Researchers asked 17,000 people:

Have you ever, when believing yourself to be completely awake, had a vivid impression of seeing or being touched by a living being or inanimate object, or of hearing a voice; which impression, so far as you could discover, was not due to any external physical cause?

When obvious cases of illness and dreaming were ruled out, 1,684 (almost 10%) said they had, and thousands of cases were published. Women reported more hallucinations than men, and visual hallucinations were most common, especially visions of a living person. Among the many hallucinations of named people, far more than could be expected by chance occurred within 12 hours either side of that person’s death. It seemed to be evidence ‘that the mind of one human being has affected the mind of another, without speech uttered, or word written, or sign made; – has affected it, that is to say, by other means than through the recognised channels of sense’ (Gurney et al., 1886, p. xxxv). Fifty years later, the psychological researcher Donald West (1948) found similar prevalence results, but unlike the original SPR survey, he found no convincing evidence for telepathy.

In the 1980s, the Launay-Slade Hallucination Scale was developed, and several surveys found large numbers of healthy people reporting experiences usually

associated with pathology, such as hearing a voice speaking one's thoughts aloud. Scores on the scale were approximately normally distributed (Slade and Bentall, 1988). Later studies revealed three factors: 1) vivid or intrusive mental events, 2) hallucinations with a religious theme, and 3) auditory and visual hallucinations. In the first case, agency is attributed to oneself and the experience recognised as one's own (*my daydreams*), whereas in 2) and 3), the experience is attributed to a source other than oneself, with salient social and agent-like properties (see also Alderson-Day and Fernyhough, 2016) sometimes extending to supernatural forces ('a voice', 'voice of God') (Waters et al., 2003). This scale has also been used to explore the rather complex relationships between the tendency to hallucinate and other variables such as reality monitoring, vividness of imagery, schizotypal personality, and susceptibility to hypnosis.

A more recent cross-cultural estimate of hallucination based on surveys from eighteen countries found that 5.2% of respondents had experienced a hallucination in their lifetime (compared to only 1.3% reporting delusional experiences involving paranoid beliefs about mind control, being followed, etc.), with lower instances in low-income countries and amongst men (McGrath et al., 2015). All this suggests that the tendency to hallucinate varies along a continuum, with pathological cases at one end, people who never hallucinate at the other, and most of us in between.

CONTEXTS AND CONTENTS OF HALLUCINATIONS

Hallucinations which fit Slade and Bentall's criteria are frequently associated with mental illness. In psychiatric conditions, including schizophrenia, bipolar disorder, and depression, around 15% report visual hallucinations and 28% auditory hallucinations. Rates are highest in schizophrenia, averaging around 27% and 59% respectively (Waters et al., 2014). Schizophrenia affects something like 0.3% of the world's population, and is difficult to define and understand; it tends to be diagnosed differently at different times and in different countries. Although the symptoms are highly variable, the core is a loss of the sense of personal control. People with schizophrenia may be convinced that other people with psychic powers are forcing their actions, or that an evil entity is controlling them. The most common kind of hallucination (reported on average by around 60% of sufferers) is hearing voices, such as aliens plotting evil deeds, or fairies chattering in the walls. Some people with schizophrenia feel that other people are inserting thoughts into their mind; some hear their own thoughts being spoken out loud as though by someone else. At their strongest, these hallucinations are detailed and compelling, uncontrollable, and experienced as completely real (Frith, 2015).

Hallucinations are also sometimes experienced as part of the 'aura' which precedes a full-blown epileptic seizure. These may be visions, disturbing smells or sounds, an intense feeling of *déjà-vu*, or even repeated scenes from memory or imagination. Patterns often develop in these experiences, which may be useful as a warning of an impending seizure and as a clue to what its triggers are or where in the brain it begins. People with dementia may also have hallucinations, with the type of hallucination depending on whether auditory or visual cortex, for example, is less affected as their condition deteriorates.

Other common causes of hallucinations are drugs, physical illness, starvation, and sleep deprivation, as well as ritual practices such as rhythmic drumming, whirling, dancing, chanting, flagellation, or control of the breath. Sensory deprivation is a powerful way to induce hallucinations. It is as though when deprived of input, our sensory systems find patterns in what little information they have, lower their criteria for what to accept as real, or turn to internally generated stimulation instead. This is simply an intensified version of the universal human habit of pareidolia: seeing familiar patterns on the flimsiest pretext, like turning lunar contours into the man in the moon or hearing messages in music played backwards.

And this is if you look carefully at certain walls soiled by different stains or at stones of uneven composition. Should you have to invent a setting, you will be able to see in these the likeness of different regions, embellished with mountains, rivers, rocks, trees, wide plains, valleys and hills in different ways; moreover, you will be able to see in them various battles and actions ready to be performed, involving strange figures, outlines of faces and clothes and endless things, which you can reduce to complete and proper shapes; in such walls or stones the same happens as with the sound of bells, in whose strokes you will find any name or word you can imagine.

(Leonardo da Vinci, *A Treatise on Painting* [Trattato della pittura], 1651, our translation)



PROFILE 14.1

Ronald K. Siegel (b. 1943)



Ronald Siegel is a pioneer of drug studies and explorer of altered states of consciousness. In the 1970s, he and his colleagues trained people to become ‘psychonauts’ — that is,

to go into altered states and report what they experienced as it happened. He has researched the effects of LSD, THC, marijuana, MDMA, mescaline, psilocybin, and ketamine, among other drugs, and has acted as consultant on several investigations of drug use. He is not just an experimenter and theoretician of psychopharmacology, but has trained in martial arts, experienced sleep paralysis, taken part in shamanic rituals, and was once locked in a cage for more than three days without food or water, all in the interests of investigating consciousness. He has a PhD in psychology from Dalhousie University and is author of many books on drugs, hallucinations, intoxication, and paranoia.

In the 1930s, British neurologist Hughlings Jackson suggested the ‘perceptual release’ theory of hallucinations: that memories and internally generated images are normally inhibited by input from the senses, and are released when that input is disrupted or absent. Louis West (1962) developed this theory, suggesting that hallucinations occur when there is both impaired sensory input and sufficient arousal to permit awareness. American psychologist Ronald Siegel (1977) likens this to a man looking out of the window near sunset. At first, all he sees is the world outside. Then, as darkness falls, the reflection of the fire inside and the room it illuminates take over, and now he sees them as though they lie outside. In this way, ‘inner’ images come to seem real.

The implication here is that either the outside or the inside of the room takes over as the current model of reality; the two compete and both cannot seem real at once. This idea has been applied to some sleep-related phenomena and to out-of-body experiences in which a completely hallucinated world takes over from the perceived world and becomes the current model of reality (Blackmore, 2009, 2017; Metzinger 2009). Something like this happens to people who immerse themselves in sensory deprivation tanks, floating in warm water in complete darkness and silence. In this situation, there is no reliable sensory input and so the self-generated world is the only reality available.



FIGURE 14.3 • The four form constants are found in decorations and works of art all over the world. Here spirals and lattices form part of a Peruvian textile design. The anthropomorphised plants in the lower corners are a cactus that produces hallucinogenic sap used for inducing visions.

Far more common, however, are hallucinations that are combined with the perceived world, as happens to lone explorers and climbers who see or hear imagined companions, and to people who become blind through either retinal or brain damage. These 'visual release hallucinations' attributable to partial blindness or as part of the adaptation to severe blindness are known as Charles Bonnet syndrome and are very common in older people who have cataracts, macular degeneration, or retinal damage through diabetes. The images are usually clear and well-defined, often featuring 'lilliputian' characters, animals, or rows of objects that are smaller than usual, and rarely threatening or scary. Sometimes sufferers don't tell anyone about their experiences for fear that they are going mad, and they can be much reassured by knowing that their situation is common (Ramachandran and Blakeslee, 1998; Menon, Menon, and Dutton, 2003).

A similar phenomenon happens with encroaching deafness when people may hear hallucinated sounds, such as hymns and ballads, choirs singing, or even whole orchestras playing. Others hear meaningless melodies, rumbling noises, or isolated words and phrases. Occasionally the sounds can be so realistic that the deaf person tries to find the source and stop them. People with tinnitus, whether temporary or long-lasting, may also sometimes hear music or voices in addition to the usual ringing sounds or white noise.

Auditory hallucinations are one of the main symptoms of schizophrenia, and here it is rare to have visual hallucinations without auditory ones too. Usually, however,





FIGURE 14.4 • (a) Hallucinated tunnels can be simple dark spaces leading to a bright light, schematic tunnel patterns, or realistic tunnels like sewers, subways, or caverns. In his experiments with THC, psilocybin, LSD, and mescaline, Siegel (1977) found that after 90 to 120 minutes colours shifted to red, orange, and yellow, pulsating movements became explosive and rotational, and most forms were lattice-tunnels, such as this one which has complex memory images at the periphery (Siegel, 1977, p. 137). (b) An almost identical tunnel was painted by David Howard, a man with narcolepsy who claimed to have been frequently abducted by aliens (see Blackmore, 2017, p. 217).

the two types occur at different times, for example a day apart. When they are fused, the two senses are typically unrelated (for example, seeing the devil while hearing the voice of a relative), suggesting independent though overlapping mechanisms.

Other hallucinations are of bodily sensations like pain, heaviness, stretching, palpitation, or touch, temperature, or proprioception (Kathirvel and Mortimer, 2013). A common tactile hallucination is formication – the feeling of ants crawling over or under the skin. These kinds of hallucinations are much rarer, harder to verify (especially in visceral cases where a doctor cannot easily look inside the body to determine that nothing is there), and not well researched. But corresponding activity in somatosensory cortices seems present as for visual and auditory hallucinations.

Phantosmia is the hallucination of a smell, usually an unpleasant smell as of something burning or rotten. This sometimes precedes epileptic seizures, and sometimes occurs in only one nostril – the one with the worse olfactory ability. It may have peripheral or neural causes, and is associated with increased activity in contralateral frontal, insula, and temporal regions which reduces after treatment of the nasal cavity.

Although there is no limit to the variety of hallucinations, there are some remarkably common features, suggesting a consistency that reflects underlying brain functions. Persistent visual forms include spirals, concentric patterns, wavy lines, and bright colours. Meditators, who may sit for long periods in front of a blank wall, report bright, variously coloured star-like bursts of light, wavelike, or cobwebby patterns, as well as shimmering, pixellation, and general brightening (Lindahl et al., 2014; Brasington, 2015). Mandalas based on circular forms are common, especially in meditative traditions, and Carl Jung included the mandala as one of the archetypal forms of the collective unconscious, describing it as the symbol of a harmonious self. These persistent patterns can be seen on shamans' drums, cave paintings, ritual designs, and clothing and artefacts from many cultures. But why?

The reason for these similarities was first investigated in 1926 by Heinrich Klüver at the University of Chicago while studying the effects of mescaline. He found that the brightly coloured images the drug induced persisted with eyes open or closed, and tended to take on four repeated forms. These 'form constants' were 1) gratings and lattices, 2) tunnels, funnels, and cones, 3) spirals, and 4) cobwebs. All are found in the hallucinations caused by drugs, fever, migraine, epilepsy, and near-death experiences, and in hypnagogic imagery and the imagery of synaesthetes.

The reason may lie in the way the visual system is organised, and in particular the mapping between patterns on the retina and the columnar organisation of primary visual cortex (Cowan, 1982; Bressloff et al., 2002). This mapping is well known from both monkey and human studies, and is such that concentric circles on the retina are mapped into parallel lines in visual cortex. Spirals, tunnels, lattices, and cobwebs map onto lines in different directions. This means that if activity spreads in straight lines within visual cortex, the experience is equivalent to looking at actual rings or circles.




FIGURE 14.5 • The mapping from eye to cortex. The visual field shown on the right is mapped onto the corresponding cortical pattern on the left. Stripes of activity in cortex are therefore experienced as though due to concentric rings in the visual field. Depending on the direction of the waves of cortical activity, either concentric rings or spirals are experienced. According to Cowan, this can explain the origin of the four form constants (1982, p. 1062).

'A hallucination is a species of reality, as capable of teaching you as a videotape about Kilimanjaro or anything else that falls through your life'

(McKenna, 1992, quoted in Rowlandson, 2012, p. 53)

One possible cause of straight lines of activation in visual cortex is disinhibition. Hallucinogenic drugs, lack of oxygen, sensory deprivation, and certain disease states can all affect inhibitory cells more than excitatory ones, causing an excess of activity which can spread linearly. The result is hallucinations of the four familiar form constants.

'Rows of mugs fixed on a wall (three rows of four) for up to two minutes. Large mugs in the top row and cups at the bottom.'

(ffychte and Howard, 1999, p. 1250)

There are also similarities in the movement, colour, and shapes of visual hallucinations. Ronald Siegel and Murray Jarvik (1975) trained volunteers to report on their hallucinations when taking a variety of drugs, including LSD, psilocybin, THC (from cannabis), and various control drugs and placebos. When the trained 'psychonauts' were given amphetamines and barbiturates, they reported only black and white forms moving about randomly, but the hallucinogens produced tunnels, lattices and webs, explosive and rotating patterns, and bright colours, especially reds, oranges, and yellows.

As for more complex visual hallucinations, they vary much more widely than the simple forms, but there are common themes too, including cartoon-like characters, scenes from childhood memory, animals and mythical creatures, fantastic cities and buildings, and beautiful scenery. In Siegel and Jarvik's drug studies, simple hallucinations came first, then a shift to tunnels and lattices, and finally more complex hallucinations. During the peak hallucinatory periods, the participants often described themselves as becoming one with the images. They stopped using similes and described their images as real. We might say that they were no longer having pseudo-hallucinations.

Visual hallucinations in degenerative eye disease include not only the form constants and vivid colours, but also visions of children, animals, buildings or landscapes, distorted faces with prominent eyes and teeth, and even copies of the same object arranged in rows or columns (ffytche and Howard, 1999). In an fMRI study, several patients with Charles Bonnet syndrome were asked to report the beginning and end

• SECTION FIVE: BORDERLANDS

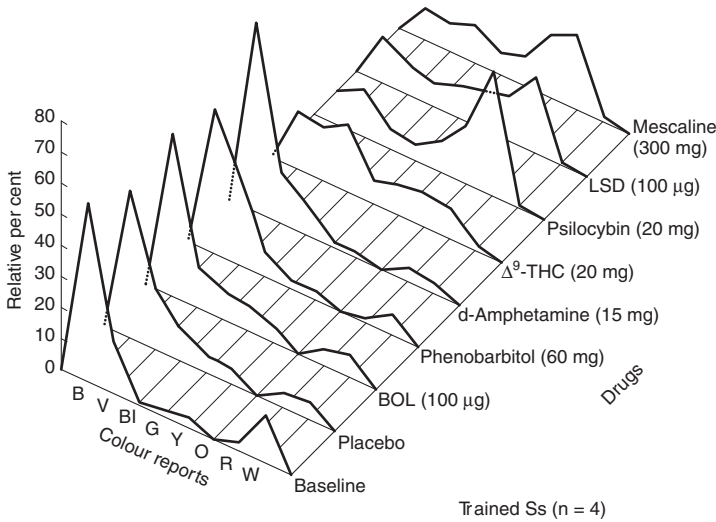


FIGURE 14.6 • Siegel and Jarvik (1975) trained psychonauts to report on their experiences while taking various drugs. Mean percentage distributions for colour are shown here. B black, V violet, Bl blue, G green, Y yellow, O orange, R red, W white.

of hallucinations while their brain activity was recorded. Hallucinations of faces were associated with activity in the face area, objects with activity in the object area, colour with colour areas, and so on. For complex visions, the features simply added up: activity in both object and colour areas was associated with a coloured object, while activity in a texture area without activity in a colour area was associated with a colourless texture (ffytche et al., 1998).

HALLUCINATING MACHINES

Hallucination-like patterns have even been replicated without human involvement. A science–arts collaboration called the Einstein’s Brain Project was inspired by the phenomena of closed-eye hallucination, including random light–dark regions, blobs, flashes, and colours in motion, as well the familiar form constants. In a camera-based experiment, the lens is covered with a uniformly illuminated goggle and bathed in yellow light to make a ganzfeld or uniform field. The video stream is then sent to a computer to be analysed for optical features, and tiny inconsistencies are tracked, amplified, and projected onto a wall. Patterns emerge from the noise as video frames accumulate and merge – just as in ganzfeld involving human participants. The authors describe the machine memory as ‘generating form from within. [. . .] It is as if algorithmic access to an archive – machine memory if you will – is, and must be, fundamentally hallucinatory (Dunning and Woodrow, 2010).

Far more dramatic images are produced by Google’s ‘deep dreaming’ algorithms. The idea is based on artificial neural networks that are trained to recognise objects in complex images. These multilayer networks are shown thousands of images and trained to extract progressively higher- and higher-level features until the final layer can identify specific objects such as faces, houses, and animals or even a specific person, breed of dog, or type of farm building. Even relatively simple networks are found to over-interpret images, finding shapes and objects that are not really there, as with human pareidolia.

The trick that researchers at Google and elsewhere have been exploring is to reverse the flow of information through the network in a process they call ‘inceptionism’, a name based on a line in the science fiction film *Inception*: ‘We need to go deeper’ (Hayes, 2015). Once the network has been trained to recognise an object (e.g. a face), the learning process is stopped and the network is run in reverse, and then the forward–backward cycle is repeated. But instead of adjusting the synaptic weights in the network, in this case the weights are held constant and the *image* (the input) is manipulated. Whether the image has the target object in it or not, as the iterations continue the self-reinforcing process produces first ghostly versions of the object and then extraordinarily complex images with multiple objects and patterns: dogs with ten legs, human heads on artificial bodies, cities sprouting snakes and eyes. They look for all the world like the psychedelic art inspired by the major hallucinogens.

Perhaps this reveals the similarity between the visual systems of animals and neural networks when they try to make sense of the world, elaborating on common features, using bits of what they know, and filling in the blanks. This all suggests that the capacity for hallucinating is an intrinsic feature of complex visual systems and is one more piece of evidence to make us wonder why we banish human hallucinations to the realm of the unreal.

HALLUCINATIONS AND THEORIES OF CONSCIOUSNESS

Does the hard problem seem any worse to you when you think about a hallucinated golden tunnel of light rather than an actual yellow-painted underground walkway with sun shining in at the far end? It should not do so. In essence, the problem is the same: as Chalmers puts it, how can physical processes in the brain give rise to subjective experience? Yet perhaps the familiarity of thinking about perceiving the 'real' world blinds us to the seriousness of the problem, which may seem more obvious when thinking about hallucinations. We know (at least roughly) what sort of cortical activity causes someone to have a potent hallucination of a bright golden tunnel. But how can the experience of a yellow tunnel (that throbbing, pulsating, realistic tunnel sucking me right now into its golden light) be caused by, or simply be, that neural activity?

For some theories of consciousness, hallucinations provide a special stumbling block. For example, sensorimotor theories entail no pictorial images or representations inside the head; instead, perceiving means having mastery of the sensorimotor contingencies between sensory input and motor responses such as moving your head, blinking, or running your fingers over something to change the input. This makes imagery and hallucinations a problem because moving, blinking, or touching them has no effect. O'Regan and Noë (2001) try to solve this problem by suggesting that knowledge of the contingencies involved is sufficient for experiencing hallucinations and imagery. In addition, the lack of feedback explains why imagery and hallucinations are not as accurate or detailed as direct perceptions.

Other theories use hallucinations in support. Higher-order theories take them as evidence that one can have a second-order thought (i.e. one can represent to oneself) that one is in a state when this is not true. For example, in the visual release hallucinations of Charles Bonnet syndrome there is higher-order representation (for example of a group of little laughing faces) without first-order representation; so higher-order representation seems to suffice for conscious




FIGURE 14.7 • Compare these two images. One is psychedelic art, the other is produced by Google's DeepDream program.

visual perception. But then we have to ask: what makes the higher-order state conscious? Why is there any what-it's-likeness to those laughing faces? And we cannot invoke a further, third-order state, for people aren't conscious of being conscious of a hallucination – they are conscious of the experience itself (Pretzman, 2012). One riposte is simply that 'the higher-order state happens to be the right kind of awareness' – the kind that we call phenomenal consciousness (Brown, 2012). But this still leaves us asking 'why?'

Dan Dennett begins *Consciousness Explained* (1991) with 'Prelude: How are hallucinations possible?', intending this to prepare the ground for his multiple drafts theory of consciousness. He proposes what would now be called a predictive-processing account, based on 'generate-and-test' theories of perception: perceptual hypotheses based on expectations and interests are constantly created and either confirmed or disconfirmed by the sensory input. This cyclical process of generate-and-test produces a model of the world that is constantly being updated but relies on having sufficient sensory input. When deprived of meaningful input, the data-driven part of the hypothesis-generating system lowers its threshold for noise. This means the answers coming back from the testing-and-confirmation part make little sense, and it goes into a random cycle of confirmation and disconfirmation. The result is hallucinations based on what the system already knows about, whether that is the simplest of geometric designs or highly detailed hallucinations produced by anxious expectation followed by chance confirmation.

This account fits with much of what we already know about hallucinations: that they are common during sensory deprivation, are induced by drugs that increase noise through cortical disinhibition and other effects, and are often elaborated into complex forms from simple beginnings.

How does this help with Dennett's theory of consciousness? If hallucination is a phenomenon of prediction and interpretation, the key point is that 'the only work the brain must do is whatever it takes to *assuage epistemic hunger*' (1991, p. 16). Sensory systems are seen not as providing a picture or representation of the world that 'enters consciousness' or is watched by the audience in the Cartesian theatre, but as continually asking multiple questions, checking against the input, and acting on the responses. This implies that a principled reality/imagination distinction is not required.

Predictive-processing accounts have made further progress with understanding hallucinations. One suggestion is that psychosis involves a breakdown in normal predictive processing. For example, people with schizophrenia fail to see many common illusions and are not susceptible to the McGurk effect, and in binocular rivalry their rate of switching between the two images is much slower than average, all implying a failure of predictive processing (Wilkinson, 2014).

More generally, in this way of understanding perception, 'your conscious percept is determined by the overall hypothesis that your brain has adopted in order to minimise prediction error' (Wilkinson, 2014, p. 148). In this way, predictive-processing accounts do away with any mystery about where the content of hallucinations comes from, and obliterate any strict dividing lines between perceptions, illusions, and hallucinations; they are all phenomena in which the brain selects the

'perceptions are hypotheses [...] – like the predictive hypotheses of science'

(Gregory, 1966/1997, p. 10)

hypothesis that best minimises prediction error. This is why our mental worlds are full of time-travel, imaginings, and dreams – as well as why we hallucinate (Clark, 2015).

So is the hallucinated tunnel ‘real’? In one sense, it is not real because there is no physically detectable tunnel present, and other people in the vicinity would not see any tunnel. In another sense, it is real because there is physically measurable activity in the person’s brain. We might also say it is real because it has measurable later effects on the person’s behaviour. This is true whether you are seeing an actual tunnel as a tunnel (vision), seeing a set of concentric circles as a tunnel (illusion), or seeing nothing-in-particular as a tunnel (imagining or hallucinating). Also, tunnels and other forms are common in hallucinatory experiences and to that extent can be shared and publicly verified. But what sort of reality is this? Should we think of any of these versions of the tunnel as ‘more real’ than any other?

‘your conscious percept is determined by the overall hypothesis that your brain has adopted in order to minimise prediction error’

(Wilkinson, 2014, p. 148)

‘A hallucination is a fact, not an error; what is erroneous is a judgment based upon it’

(Russell, 1914, p. 173)

EXTRA-SENSORY PERCEPTION

Even if hallucinations on the strictest definition (not knowing you are hallucinating) are relatively rare, there is no doubt that hallucinations exist: their existence *is* the experience. Other phenomena on the borders of the imagination make stronger claims for the nature of their reality – for example, that what looks like ‘mere imagining’ may be a form of mental travel or communication at a distance. And this takes us into the realm of parapsychology.

Levels of belief in the paranormal are high (Gallup and Newport, 1991; Blackmore, 1997), and if telepathy, precognition, or any other paranormal phenomenon did occur, this would have truly extraordinary implications for how we understand the universe, and perhaps for the science of consciousness in particular. Although it is not entirely logical, psychic phenomena are popularly thought to be evidence for the ‘power of consciousness’, due to ‘consciousness interactions’ or ‘consciousness-related anomalies’. Proof of their existence is sought in the hope of overthrowing materialist theories of mind and demonstrating that consciousness is independent of time and space. American parapsychologist Dean Radin argues in *The Conscious Universe* that ‘Understanding [paranormal] experiences requires an expanded view of human consciousness’ (1997, p. 2). Cardiologist Pim van Lommel (2013) claims that near-death experiences are evidence for ‘non-local consciousness’ and even ‘endless consciousness’.

There probably are no paranormal phenomena (Blackmore, 1998; compare Bem, 2011 and Galak et al., 2012), and if not, the widespread beliefs and frequent reports of psychic experiences must be explained some other way. Pinpointing the position they occupy on the spectrum of reality and imagination is one way of doing so.

Parapsychology was the brainchild of J. B. and Louisa Rhine, two biologists at Duke University in North Carolina (Mauskopf and McVaugh, 1980). Like the British psychological researchers before them, they had far-reaching ambitions. They wanted to find evidence against a purely materialist view of human nature, and fight against the powerful behaviourism of their time. They thought that their new science was the way to demonstrate the independent agency of the mind, and even to solve

the mind–body problem. The Rhine Research Center, which still carries on their work, defines parapsychology as ‘the scientific study of interactions between living organisms and their external environment that seem to transcend the known physical laws of nature. Parapsychology is a component of the broader study of consciousness and the mind’.

The Rhines began by defining and operationalising their terms, with J.B. Rhine’s first book, in 1934, launching the term ‘extra-sensory perception’. ‘ESP’ was proposed as a general term to include three types of communication without the use of the senses: telepathy, in which the information comes from another person; clairvoyance, in which it comes from distant objects or events; and precognition, in which it comes from the future. In addition, the term ‘psi’ covers both ESP and psychokinesis (PK), that is, the effect of mind over matter or the ability to influence things at a distance without any physical interaction. These terms are still defined this way in parapsychology, although their popular meanings are rather different. (Our website has more material on parapsychology, including ESP, PK, and some of the many controversies in the field.)



PRACTICE 14.1

LIVING WITHOUT PSI

The possibility of ESP is comforting. We might sense when a loved one is in danger, share our deepest feelings with others, or find ourselves guided by a supernatural power. For this exercise, try living without such comfort.

If you believe in psi, or angels, or life after death, or spirits, take this opportunity to live without them. If you catch yourself imagining a helping spirit or guardian angel, watch what comes to mind and gently let the image go. If you find yourself imagining that someone you know who has died is still around, watching you or caring what you do, observe where your imagination leads you and the effects it has on you. You need not abandon your beliefs forever. Just set them aside for a few days and see how the world looks when you know you are completely on your own.

Sceptics should do this too. You may be surprised to find yourself *willing* something to happen even though you know you cannot affect it, or conjuring up an image of a friend hoping they will know when you need them. Ask yourself this. Do we live better or worse for a belief in the supernatural? Don’t give a glib, intellectual answer. Look and see what happens when you try to root it out completely

‘remote viewing must signify the existence of an astonishing hidden human potential’

(Targ and Puthoff, 1977, p. 9)

For testing ESP, Rhine used a special pack of twenty-five ESP or Zener cards, with five distinctive designs: square, circle, star, cross, and wavy lines. For telepathy, a receiver, or percipient, had to guess the order of a pack of cards being looked at by a sender, or agent. For testing clairvoyance, the cards were shuffled and hidden from view; for precognition they were shuffled only after the receiver had made the guesses. Rhine reported many successful results but not without much controversy

and many failures of replication (Irwin and Watt, 2007). In general, this method of ‘forced-choice’ guessing with boring cards obtained only extremely weak effects – if indeed they were effects at all. For this reason, by the 1970s various ‘free-response’ methods were developed which, although more time-consuming, are much more enjoyable to do.

In ‘remote viewing’, for example, a target person goes to a randomly selected remote location and looks around for a specified length of time. Meanwhile, the receiver sits and relaxes, reporting any impressions or images that arise.

Afterwards, either the receiver or an independent judge tries to match up the impressions with a limited set of possible target locations and pick the right one. This means that although the descriptions are given freely, inferential statistics can be used to test the results. Remote viewing became famous when physicists Russell Targ and Harold Puthoff (1977), at the Stanford Research Institute in California, obtained highly significant results. Then two psychologists, David Marks and Richard Kammann, argued that there were clues in the transcripts which might have been used to obtain spurious results. This led to a controversy in the prestigious journal *Nature* and attempts by others to determine the relevance of these clues (Marks, 2000).

In 1995, the American Institutes for Research reported on ‘Stargate’, a 24-year, \$22-million government-funded research project on the feasibility of using psychic powers for intelligence gathering. Many of their experiments used the same remote-viewing protocols, but arguments about the adequacy of the methods used and the significance of the results followed (Hyman, 1995; Utts, 1995; Wiseman and Milton, 1998). American statistician Jessica Utts described Stargate as providing some of the most solid evidence of psi to date, whereas Marks described it as ‘a series of closed-off, flawed, nonvalidated, and nonreplicated studies’, concluding that ‘Remote viewing is nothing more than a self-fulfilling subjective delusion’ (Marks, 2000, p. 92) – that is, the remote-viewers or the experimenters imagine connections with the target, even though if the protocols are adhered to correctly, their imaginings should have no effect. Regardless of who is right, the US government decided that remote viewing could not be used for gathering intelligence. Documents from the Stargate project were released and made available online in January 2017. There is no evidence that any other country has successfully employed ESP for spying.

Targ continued his research and used the findings of remote viewing to illustrate his contention that most of us have untapped psychic powers. The clutter of our conscious minds suppresses our natural abilities, he claims, but if we learn to quiet this noise through meditation and other forms of self-inquiry, we may be able to experience what the remote-viewing data show: ‘without a doubt, that our mind is limitless and that our awareness both fills and transcends our ordinary understanding of space and time’ (Targ, 2004, p. xiii; see Activity 14.1).

Even more controversy ensued over another method for testing ESP, this time in the ganzfeld (German for *total field*). Participants in a ganzfeld experiment lie comfortably, listening to white noise or seashore sounds through headphones, and

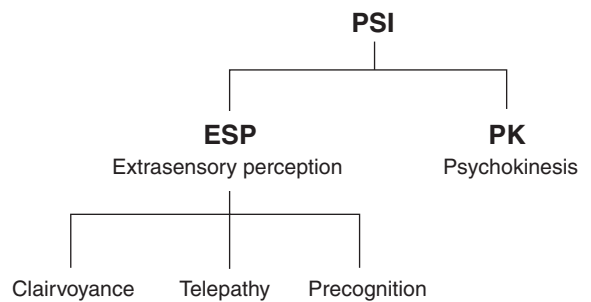


FIGURE 14.8 • Terms used in parapsychology. ‘Psi’ is a general term that refers to all kinds of paranormal phenomena or the supposed mechanism underlying them. There are four forms of psi and three types of ESP.



ACTIVITY 14.1

Telepathy tests

1 A (reasonably) controlled experiment

The problem with testing for telepathy is the many ways in which subtle, but normal, communication can appear to be telepathic. In experiments with cards or pictures, for example, there is not only the possibility of sensory leakage via subtle sounds, movements, or deliberate fraud, but if the targets to be guessed at are not properly randomised, people's natural tendency to prefer certain targets or even orders of targets can produce spurious results. Even worse, if the 'sender' and 'receiver' know each other and can choose the target, they are likely to choose the same things. Bearing this in mind, it can be fun to try experiments which allow these faults before comparing it with one that does not. Here is a reasonably controlled experiment that can be done in class (a sample answer sheet, as well as a more basic experiment and an impressive demonstration, can be found on the website).

Advance preparation. Remove the court cards from a pack of playing cards, leaving forty cards of four suits. Use a random number generator (not shuffling) to decide the target order. Assign 1 – hearts, 2 – spades, 3 – clubs, 4 – diamonds. Make a record of the target order. Arrange the cards in that order with the first card on the top when the pack is face down. Place an unused card on the bottom to conceal the last card. Seal the pack in an opaque envelope and the list in another envelope. Get two stopwatches, prepare answer sheets as provided on the website, and find a room for the sender.

The experiment. Choose someone to be sender, give her a watch and the sealed pack, and arrange the exact time at which she will turn over the first card. She then goes to the appointed room, opens the envelope, and places the pack face down on the table. At the pre-arranged time she turns over the first card and concentrates on it, turning over the rest at 15-second intervals. The whole test will take 10 minutes. Meanwhile, you call out the numbers 1–40

wear half ping-pong balls over their eyes to produce a uniform white or pink field – the ganzfeld. This tends to produce a very relaxed state with free-flowing imagery, and researchers hoped this would be conducive to psi success. While in the ganzfeld, people report what they experience, and this is recorded for judging afterwards. Meanwhile, a sender in a distant room views a picture or video clip, the target. After half an hour or so, participants are shown four such pictures or videos, and are asked to decide which one was the target, so that as with remote viewing, a free-response method culminates in a forced choice. Alternatively, independent judges make the decision by comparing the transcripts of the session with the possible targets. Claims of success and counter-claims of failures led to the 'Great Ganzfeld Debate' (Concept 14.1).

A more everyday blurring between imagination and reality is where people claim they can feel when someone is staring at them, or conversely can get people to turn around by staring at them. British researcher Rupert Sheldrake is among those who have experimented on 'remote staring' using video cameras or online methods to avoid sensory leakage and claiming highly successful results. These, he argues, suggest the existence of a perceptual field that extends beyond the head, and are 'more compatible with theories of vision that involve both inward and outward movements of influence' (2005, p. 32). In other words, some paranormal influence extends from the eye of an observer and is detectable by the observed. His numerous critics have accused him of many inconsistencies, mistaking signal for noise, over-hyping weak evidence, and being deeply confused over theories of vision (2005, with commentaries).

Let us suppose that experiments like these one day produce reliable evidence for ESP. What would the implications be for understanding consciousness? Interestingly, although many researchers claim that this would prove the power of consciousness, or the independence of mind, there is little in the experiments to support this claim. Even in the most successful ESP experiments, participants are not consciously aware of which guesses are hits and which are misses. If they were, these guesses could be separated out and the scoring rate dramatically improved, and this has never proved possible.

Some methods, such as the ganzfeld, do involve a mild 'altered state' of consciousness, but there is no evidence that people who enter a 'deeper' state do better, or that an altered state is necessary for success in the ganzfeld. There is also no consensus over what it is about the ganzfeld that makes it psi-conducive, if indeed it is, other than that it provides a blank slate on which the imagination can trace its images.

Hypnosis has also been used as an induction technique, but again there has been no clear demonstration that, even if

something is qualitatively altered in hypnosis (see Chapter 13), this change in consciousness helps. Many more experiments would be needed to establish which of a potentially unlimited number of untested variables are relevant. The connection with consciousness seems to come entirely from theoretical suppositions about how ESP might work. There is no direct evidence that consciousness is involved in any way (Blackmore, 1998). We have considered evidence in earlier chapters that the influence of conscious will on our own actions, attention, and perception, let alone anyone else's, may be illusory. So maybe we should expect no such involvement, whether in terms of the power of conscious will or the power of consciousness perception beyond the material.

CONJURING OTHER WORLDS

All of us conjured other worlds when playing as children: inventing food and drinks for dollies' tea-time, imagining illnesses to be cured by 'doctors and nurses', and creating invisible cargoes to be carried by toy trucks on imaginary roads. Many children, especially only children, have imaginary playmates. Some children play and talk with the same friend for many years, though not often past the age of ten. In the early years, the playmates are described as solid and real, but older children rarely see them that way. Most imaginary companions are people, usually of the same sex as the child, but they can be animals, invisible toys, storybook characters, and even things like clouds or door-knobs (Siegel, 1992). These friends take part in conversations, games, and all sorts of creative activities.

Pretend play is crucial to how children develop their causal understanding of the physical and mental world. In one series of experiments, two-year-olds watched Naughty Teddy victimising other toy animals with make-believe substances: for example, squirting pretend toothpaste on to a rabbit's ear (Harris, 2000, pp. 17–19). When asked to describe what they saw, the children referred to the pretend substances and actions ('toothpaste' and 'squeezed') and the consequences of those actions (making the rabbit's ear 'dirty' or 'wet') despite having the objectively dry and clean rabbit in front of them. Then a brick on a paper plate is pushed towards the rabbit and children are told the rabbit likes eating banana. This time Naughty Teddy squeezes toothpaste onto the brick instead of the rabbit, and the children didn't just infer a non-existent substance (toothpaste) but also swapped the real object's name (brick) for the make-believe name (banana); if they failed to do this, they almost always just pointed to the brick or remained silent. This suggests that they knew the causal outcome was directed at the object the prop stands for, not the prop itself. These imaginary worlds are robust, but they also obey the same causal laws as the real world

at the correct times and the receivers write down which suit they think the sender is looking at.

When the test is complete, ask the sender to return. Call out the target sequence and ask each person to check their neighbour's scores. If you have a large enough group (say twenty or more), you can show the results by building up a histogram for all to see. Ask each person in turn to say how many hits they got, and add each result to the growing picture. At first the results may seem impressive, or strange, but they will tend ever closer to a normal distribution with a mean at 10. If the results deviate from 10 and you wish to test them statistically, use a normal approximation to binomial, or a one-sample t-test using 10 as the expected value (but see below).




FIGURE 14.9 • Sample histogram of results from a simple ESP experiment. Unless psi is operating in your experiment you are likely to get a normal distribution with a mean of 10. As you gather more data, the histogram approaches more closely a normal distribution.

This method avoids most of the obvious problems, but some remain, including the 'stacking effect', which means that t-tests are not wholly appropriate when many people guess at the same target list. Sensory leakage or fraud might still have taken place, and you may discuss whether they could ever be completely ruled out. Other easy psi experiments, with more detailed instructions, can be found in Blackmore and Hart-Davis (1995).



CONCEPT 14.1

THE GANZFELD CONTROVERSY

The first ganzfeld experiment was published in 1974 by the American parapsychologist Charles Honorton. Attempts at replication produced varying results, steadily improving techniques, and many years of argument, all culminating in the 1985 'Great Ganzfeld Debate' between Honorton and American psychologist Ray Hyman (1985). Both carried out meta-analyses of all the available published results, but they came to opposite conclusions. Hyman argued that the positive results could all be explained by methodological errors and multiple analyses. Honorton argued that the overall effect size was large and did not depend on the number of flaws in the experiments, and that the results were consistent, not dependent on any one experimenter, and revealed regular features of ESP. In a 'joint communiqué' (Hyman and Honorton, 1986), they detailed their agreements and disagreements, and made recommendations for the conduct of future ganzfeld experiments.

In 1994, the original meta-analysis was republished in *Psychological Bulletin* (Bem and Honorton, 1994), along with impressive new results obtained with a fully automated ganzfeld procedure carried out at Honorton's Psychophysical Research Laboratory (PRL) in Princeton. This 'autoganzfeld' was hailed as a fraud-proof technique that would finally provide a repeatable experiment for parapsychology, but criticisms began again with the suggestion that acoustic leakage might have occurred (Wiseman, Smith, and Kornbrot, 1996).

Another problem concerned nine studies carried out by British psychologist Carl Sargent at Cambridge University. These nine comprised nearly a third of the twenty-eight studies in the original meta-analysis and had the second highest effect size after Honorton's own. Having failed to obtain significant results in her own experiments, Sue visited Sargent's laboratory in 1979 and found that the experiments, which looked so well controlled in print, were far from fraud-proof. She uncovered serious errors

does, and through this kind of play, children develop their knowledge about reality by stepping back from or going beyond it.

This capacity for creating other characters and other worlds continues into adulthood in daydream fantasies and in the enjoyment of fiction- and poetry-reading, film-viewing and theatre-going, electronic gaming, creative writing, painting, and other arts, which have been described as 'qualia machines' offering up new varieties of consciousness (Reinerth and Thon, 2016). When we feel 'immersed' or 'absorbed' in, or 'transported' to, a world created by a written text or a set of moving images, we may retain more or less awareness of the environment in which we are reading or watching. This might depend on many other factors, including our evaluations of or empathy with the protagonists, the richness of our mental imagery, our familiarity with the story's genre, and maybe even basic demographic factors like gender (van Laer et al., 2014).

People also vary widely in their capacity for 'psychological absorption', a variable closely related to hypnotisability. Absorption is usually measured with the Tellegen Absorption scale (Tellegen and Atkinson, 1974; Jamieson, 2005) and those who score highly are more likely to report a variety of unusual experiences and respond more strongly to drugs like psilocybin (Blackmore, 2017).

Virtual reality technologies can now create elaborate multisensory simulations which are heightened by users' ability to interact physically with them. VR worlds can induce motion sickness, or 'simulation sickness', thanks to how they manipulate sensory perception and feedback, and responses to VR on dimensions like social paranoia or degree of presence in the virtual world can be used to predict future occurrence of PTSD symptoms (Freeman et al., 2014). And so, the boundaries between 'consensus reality' and other kinds of reality continue to shift and blur.

In contemporary Western culture, other worlds are usually confined to shared forms of fiction, or to private fantasy, but in many other cultures they are deliberately cultivated

and shared as closer equivalents to the everyday world. In many cultures, certain people train as 'shamans'. This word came originally from the Siberian Chuckchee tribe but is now widely used to describe men and women who can enter spirit worlds, cure sickness through magic, or contact spirits and other invisible beings. Usually shamans follow elaborate rituals, often but not always involving hallucinogenic drugs, to reach these other worlds (Krippner, 2000).

One such culture is that of the Yanomamö, a group of indigenous people living deep in the forest between Venezuela and Brazil (Chagnon, 1992). Their world of myths and invisible entities consists of four parallel layers, one above the other, including the third layer of forests, rivers, and gardens in which they live. Accomplished shamans can call the beautiful *hekura* spirits from the sky, hills, trees, or even from the edge of the universe to enter their bodies through the chest and there to find another world of forests and rivers within.

To call *hekura*, the shamans (who in this culture are only ever men) prepare a complex hallucinogenic green powder called *ebene*, paint themselves elaborately with red pigment, put on their feathers, and blow the powder into each other's nostrils through a long hollow tube. Coughing, gasping, groaning, and dribbling green mucous from the nose, they then call the *hekura*, who soon come glowing out of the sky along their special trails into the shaman's chest, from where they can be sent to devour the souls of enemies, or to cure sickness in the village.

Sometimes researchers have been invited to join such ceremonies and take the drugs themselves. Siegel describes a long night spent with a Huichol Indian shaman in Mexico, matching him gulp for gulp in drinking a potent alcoholic liquor made from the agave plant and a gruel made from the peyote cactus, which contains the hallucinogen mescaline (Chapter 13). When the first waves of nausea had passed, Siegel opened his eyes and 'the stars came down', darting about and leaving tracer patterns in the air. When he tried to grab one, a rainbow of afterimages followed his moving hand. Then there were patterns, all the familiar form constants, and much more. A lizard crawled out of his vomit, followed by thousands of army ants in party hats. 'Stop it! I want answers, not cartoons!' he pleaded, and he asked the shaman about hallucinations. The answer came clear: 'There are no hallucinations with peyote. There are only truths' (Siegel, 1992, pp. 28–29).

Back home in his California laboratory, Siegel knew that what he had seen was all in his own mind.

How do you tell this holy man who believes he has the power to see the gods that there are no more gods or Demons than there are images of those things in the brain? How do you tell a poor naked farmer who has only his peyote dreams that the world of our dreams is all inside our minds?

and failures to follow the protocol. She concluded that Sargent's results, and therefore the meta-analyses that relied so heavily on them, provided no reliable evidence for psi (Blackmore, 1987; Sargent, 1987), a conclusion that spurred her transformation from belief in ESP to scepticism (Blackmore, 1996a).

Following the apparent success of the autoganzfeld, more replications followed, but few were successful. Then another meta-analysis of thirty new studies found no evidence for ESP (Milton and Wiseman, 1999), while others, including further new studies, did (Bem, Palmer, and Broughton, 2001; Williams, 2011). The arguments have continued without resolution (Milton, 1999; Palmer, 2003).

'the conceptually infused alternatives to reality that children conjure up feed back on their assessments of reality'

(Harris, 2000, p. 7)
'science fiction [is] the only genuine consciousness-expanding drug'

(Arthur C. Clarke, 'Of sand and stars', 1983)

'There are no hallucinations with peyote. There are only truths'

(Huichol shaman, in Siegel, 1992, pp. 28–29)

'How do you tell a poor naked farmer who has only his peyote dreams that the world of our dreams is all inside our minds?'

(Siegel, 1992, p. 31)

(Siegel, 1992, p. 31)

But we may wonder whether this distinction between ‘real’ and ‘in the mind’ is all that clear. The *hekura* dancing down their shimmering trails and the stars coming down from the sky are not physical, publicly measurable objects. Yet they have been seen again and again by countless peoples separated in time and space in their different cultures across the world. To this extent, they are publicly available. If you took the right mixture of drugs, in the right setting, you would see them too. What sort of reality does that give them? And what is it about the mental that makes us so hesitant to call it real?

One controversial player on this edge of reality is anthropologist Carlos Castaneda, famous for his many books about his teacher, the Yaqui Indian Juan Matus (Castaneda, 1968). As the story goes, Castaneda first met the old *brujo*, or medicine man, in the summer of 1960 at a bus depot in a border town in Arizona. While Castaneda prattled on about how much he knew about peyote, and what he wanted to learn, Don Juan peered at him patiently with shining eyes, knowing that Castaneda was talking nonsense. But they met again and Castaneda became Don Juan’s apprentice for four years. This ‘man of knowledge’ taught his disciples sorcery, taking them through strange rituals and journeys, and using three hallucinogens: peyote, which contains mescaline; jimson weed or datura, which contains tropane alkaloids including atropine; and mushrooms containing psilocybin (Chapter 13). According to Don Juan, peyote teaches the right way to live, while the other drugs are powerful allies that can be manipulated by the sorcerer. Castaneda suffered ordeals of sickness, pain, confusion, and whole worlds of visions that were, according to Don Juan, not hallucinations but concrete aspects of reality. Castaneda dubbed them ‘a separate reality’ (1971).

After many years of training, Castaneda began learning to ‘see’: a non-ordinary way of looking in which people appear as fibres of light, as luminous eggs in touch with everything else and in need of nothing. His head once turned into a crow and flew away, he heard a lizard speak, and he became a brother to the coyote. On one occasion he used jimson weed to fly, as medieval witches were said to do by using the chemically related deadly nightshade, *Atropa belladonna*. He argued with himself and with Don Juan that his actual physical body could not have flown, yet it apparently ended up half a mile from Don Juan’s house. Finally he learned to keep death ever-present and not to be so concerned with his ordinary self – indeed to stop the internal dialogue and erase his personal history.

Similar experiences are reported with ayahuasca, the hallucinogenic drink made by Amazonian shamans and used for healing, insight, and many other purposes (Chapter 13). The effects last many hours, with after-effects sometimes going on for days, and can be varied and controlled for different purposes by fine variations in the method of preparation. Most common are colourful visions of snakes and serpents, as well as bodily distortions and even the sense of being transformed into another creature or transposed into other worlds of living plants and animals.

Luis Eduardo Luna, Director of the Research Center for the Study of Psychointegrator Plants, Visionary Art, and Consciousness, in Brazil, offers a detailed description of the visual experiences induced with ayahuasca. Although the visions often move with head and eye movements as you would expect in a hallucinatory experience, sometimes ‘it is as if I am totally immersed in a three-dimensional world, so that, as in the real world, when turning my head, different things would be perceived’ (2016, p. 258). He notes that sensorimotor possibilities are limited in the visionary realm:

for example, he cannot change his perspective so as to be able to see what is behind objects in front of him. But other factors result in a heightened, not a lessened, experience of reality. Because ayahuasca visions are not mediated by the cornea, iris, and lens, the gradations in acuity that shape the experiences of normal vision (with resolution highest at the fovea and much lower at the periphery) are absent: 'everything in the inner field of vision seems to be equally sharp, which may contribute to the "more real than real" feeling that is so frequently reported in ayahuasca experiences' (p. 259).

Experienced ayahuasca users travel in this world or other worlds, according to their traditions, and describe non-ordinary ways of seeing. They claim that the gods, demons, heavens, and hells that they visit are as real as, or even more real than, the ordinary world of normal vision. They describe gaining spiritual insights and a deeper understanding of reality and of themselves.

For Luna (2016), the experiences elicited by ayahuasca overwhelm him with the feeling that much more is going on than simply the constructions of his own mind. Drinking ayahuasca makes the idea that consciousness is limited to humans seem ludicrous. 'The feeling is rather that consciousness permeates everything, that it might be primordial' (p. 268), and that sacred plants are just one way for humans to tap into its various manifestations.

Could this be?

BUT IS IT REAL?

What are we to make of all this? Like Castaneda in his sceptical anthropologist mode, we may claim that the experiences are 'all in the mind': that they are imaginary and not real. Indeed, it turns out that Castaneda's books themselves were more works of fiction than ethnographic records of research. Writer Richard de Mille made a thorough study of Castaneda's works and concluded that 'Marked anachronisms or logical conflicts in Castaneda's work must argue that his text is an imaginative fabrication rather than a factual report' (1976, p. 197). The results were a mess: 'The wisdom of the ages folded into an omelet with the neurosis of the century' (p. 18). Yet Castaneda does force us to wonder about the nature of hallucinations.

A character is either 'real' or 'imaginary'? If you think that, hypocrite lecteur, I can only smile. You do not even think of your own past as quite real; you dress it up, you gild it or blacken it, censor it, tinker with it . . . fictionalize it, in a word, and put it away on a shelf – your book, your romanced autobiography. We are all in flight from the real reality. That is a basic definition of Homo sapiens.

(John Fowles, *The French Lieutenant's Woman*, 1969/2004, p. 97)

Take those luminous eggs and radiating fibres, reminiscent of the haloes of Christian saints and the auras of the Theosophical tradition. Auras are a good example of something that is commonly reported, has consistent features, and yet is not




FIGURE 14.10 • The doorway test for auras. The psychic claimant stands facing the edge of an open doorway. A target person, whose aura the psychic says he can see clearly, takes one of two possible positions; perhaps five times each, in random order. At position (a) neither she nor her aura should be visible; at position (b) her body is not visible but her aura should easily be seen, sticking out past the side of the doorframe. On each trial the psychic must say whether he sees the aura sticking out or not. There is no published evidence that anyone has ever passed the doorway test, suggesting that whatever auras are, they are not physically present in the space around the body.



FIGURE 14.11

physically present. Kirlian photography, which is sometimes claimed to record auras, actually measures the corona discharge from charged surfaces, and Kirlian photographs do not resemble seers' descriptions of auras. And no one has ever passed the 'doorway test' designed to find out whether psychic claimants can see an aura sticking out from behind a wall (Tart, 1972b; Blackmore, 2017).

Seeing auras may seem trivial, but the lessons learned from other-world experiences should perhaps not so lightly be dismissed. Those who are experienced in the use of hallucinogenic drugs learn things that no novice has any inkling of. They learn to look calmly into their very worst fears, face up to death, confront or lose themselves, and many other lessons. Special skills are needed for exploring the worlds revealed this way, and those who acquire this kind of wisdom recognise it in others. Understanding all these phenomena is not helped by trying to find a sharp line between reality and imagination.

Some kind of distinction is needed, however, otherwise we would not be able to make judgements about the reliability of eyewitness evidence after a crime, or reassure someone who hallucinates a threatening figure that they need not be afraid. But when we say things like 'it's all in the mind' to mean that the realm of the imagination, or of the mind more generally, is *unreal*, we go too far, because body, mind, and environment are always linked. Going too far has wide-ranging and serious consequences, from denying that mental illness is really real (Chapter 13) to dismissing a whole spectrum of forms of consciousness as irrelevant to the investigation of 'consciousness itself' – whatever that is.

You may object, however, that we have chosen only the exceptional experiences of shamans, adepts, and drug users; so here is something that can happen to anyone.

I was lying on my back in bed and drifting off to sleep, when I found I couldn't move. There was a horrible buzzing, vibrating noise, and I was sure there was something – or someone – in the room with me. I tried desperately to see who it was but I couldn't move anything but my eyes. Then a hideous dark shape with an evil smell loomed up over the end of my bed and lurched towards me. I tried to scream but no noise came out. The dark shape came closer and closer and forced itself on my chest, pressing down so I could hardly breathe. It seemed to be speaking to me but I couldn't make out the words. Then it dragged on my arms and legs and began pulling me out of bed.

Imagine that this experience happened to you. What would you think as you struggled to cope? What would you think once your heart stopped pounding and the smell of the creature left your nostrils? Would you comfort yourself with the thought that the menacing black figure wasn't *real* at all and was only *imagined*? Or might you decide that it was an alien come to abduct you, or perhaps the ghost of someone who had died? Either way, you face a problem. If the creature was *real*, why did the door remain closed and the bed covers undisturbed? Why did no one else see the creature coming through the house? Obviously it wasn't real in that public sense. On the other hand, if it was only *imagined*, how could it have such a powerful effect on you, and make your heart pound and your hands

'there's nothing odd about the sense in which subjective phenomena can be objective facts [...] pain is real!'

(Strawson, 2011, p. 265)

sweat? Obviously *something* happened to you, and the experience itself was real enough, wasn't it?

In the next chapter, we will explore sleep, dreams, and some further weird experiences that haunt the borderlands of sleep, including the example of sleep paralysis just described. As with many of the 'altered states' that we considered in the previous chapter, exploring these borderlands suggests that many of the other distinctions we are so familiar with start to melt away – not just reality versus imagination, but body versus mind, self versus other, and conscious versus unconscious.

Blackmore, S. J. (2001). What can the paranormal teach us about consciousness? *Skeptical Inquirer*, 25(2), 22–27.

Uses ganzfeld research to argue against belief in the paranormal power of conscious will.

Luna, L. E. (2016). Some observations on the phenomenology of the ayahuasca experience. In L. E. Luna and S. F. White (Eds), *The ayahuasca reader: Encounters with the Amazon's sacred vine* (2nd ed.) (pp. 251–279). Santa Fe: Synergetic.

Draws on extensive personal experience to discuss visual experiences induced by ayahuasca and their relation to lucid dreaming, creativity, and reality.

Schacter, D. K., Addis, D. R., Hassabis, D., Martin, V. C., Spreng, R. N., and Szpunar, K. K. (2012). The future of memory: Remembering, imagining, and the brain. *Neuron*, 76(4).

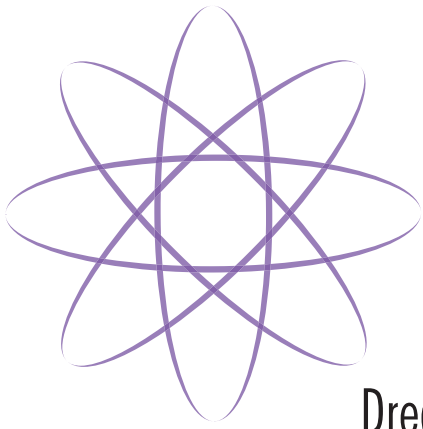
Similarities and differences between remembering the past and imagining the future.

Sheldrake, R. (2005). The sense of being stared at, Part 1: Is it real or illusory? *Journal of Consciousness Studies*, 12(6), 10–31. Peer commentaries (pp. 50–116), especially Braud, French, Radin, and Schlitz; and response, pp. 117–126.

Explores evidence for the idea that people know (without seeing) that someone is looking at them.

Siegel, R. K. (1992). The psychonaut and the shaman. In R. K. Siegel, *Fire in the brain*. New York: Penguin.

Hallucinatory adventures with Siegel's 'psychonaut team' and a Mexican shaman.



Dreaming and beyond

C
H
A
P
T
E
R

FIFTEEN

I was on a ski lift, a double-seater chair, moving slowly up into the high snowy peaks. It was cold and dark – nearly dawn, and the deep blue sky was lightening where the sun was about to break through. ‘But this lift isn’t supposed to open until 8.30 a.m.’, I thought. ‘How did I get here? Lifts don’t run in the dark. What’s going on?’ I began to panic. I looked down and realised that I had no skis on, and you need skis to get safely off the lift. I would have to run and hope I didn’t fall. As my boots were about to hit the ground I suddenly knew the answer. I was dreaming, and with that realisation it was as though I had woken up. Of course lifts don’t run in the dark. I looked around, conscious in my own dream, gazing at the beauty of the morning mountains as the sun streamed over the top.

This is an example of a lucid dream: a dream in which you know *during the dream* that you are dreaming. This ability to ‘wake up’ inside a dream while staying asleep prompts all sorts of interesting questions about sleep, dreams, and ‘altered states’ of consciousness. What does it mean to say that I ‘wake up’ or ‘become conscious’ in a lucid dream? Aren’t you conscious in ordinary dreams? What are dreams anyway? Are they experiences or only stories constructed on waking up? And who is the dreamer?

In this chapter, we will skim over the basics of sleep and dream research, for they are well covered in many texts (e.g. Empson, 2001; Hobson, 2002; Horne, 2006; Moorcroft, 2013), and concentrate on what ordinary dreams, as well as some more exotic kinds of dream and sleep-related phenomena, can tell us about consciousness. We ask the same question about out-of-body and near-death experiences,

both of which disrupt our ordinary sense of being a conscious self who looks out from inside our skin.

WAKING AND SLEEPING

I've dreamt in my life dreams that have stayed with me ever after, and changed my ideas: they've gone through and through me, like wine through water, and altered the colour of my mind. And this is one: I'm going to tell it – but take care not to smile at any part of it.

(Emily Brontë, *Wuthering Heights*, 1847)

Every day we all go through a cycle of three states: waking, REM (rapid eye movement) sleep, and non-REM sleep, a typical night's sleep consisting of four or five cycles between non-REM and REM sleep, and often some unremembered micro-awakenings. These waking and sleep states are defined by physiological and behavioural measures, including how easily the person can be awakened, their eye movements and muscle tone (the degree of passive contraction in the muscle fibres), and their brain activity as measured by either EEG or scans. In REM sleep, the brain is highly active and the EEG resembles that of waking, although paradoxically, the sleeper is harder to wake up than during non-REM sleep. Even in non-REM sleep, the overall firing rate of neurons is as high as in waking states, but the pattern is quite different, with the EEG dominated by long, slow waves rather than complex, fast ones.

The neurochemistry and physiology of these states is well researched. For example, the neuromodulators adenosine and melatonin play crucial roles in inducing sleep. During sleep, the REM cycle is controlled by the reticular formation in the pons in the brainstem, and not by higher brain areas, which are unnecessary for normal sleep cycling. Within the brainstem are cholinergic REM-on nuclei and aminergic (both noradrenaline and serotonin) REM-off nuclei, which reciprocally activate and inhibit each other and control the switching of states.

During sleep, parts of the brain are isolated in different ways and to different extents. Blocking of sensory input happens at the thalamocortical level in non-REM sleep and at the periphery in REM sleep. There are also different phases of REM, and fMRI studies show that during tonic (persistent) REM auditory stimuli still activate auditory cortex to some extent, while during phasic (intermittent) REM, when eye movements and muscle twitches occur, the brain operates in a functionally isolated closed loop (Wehrle et al., 2007).

In REM sleep, the brain stem blocks motor commands at the level of spinal motor neurons so that whatever is going on in motor cortex does not result in physical activity. This means you can dream of climbing out of the window onto the roof, but your legs won't let you do it – although these protective mechanisms can break down briefly in sleepwalking, and are overactive in sleep paralysis. At the same time, the pons, amygdala, hippocampus, and anterior cingulate are especially active, as are parts of the visual system and visual association areas, but the dorsolateral prefrontal cortex (associated with executive functions like working memory, problem-solving, and planning, as well as motor organisation) is much less active than during waking.




FIGURE 15.1 • Behavioural states in humans. States of waking, NREM sleep, and REM sleep have behavioural, polygraphic, and psychological manifestations. The sequence of these stages is represented in the polygraph channel. Sample tracings of three variables used to distinguish state are also shown: electromyogram (EMG), which is highest in waking, intermediate in NREM sleep, and lowest in REM sleep; and the electroencephalogram (EEG) and electro-oculogram (EOG), which are both activated in waking and REM sleep and inactivated in NREM sleep. Each sample is approximately 20 seconds (Hobson 2002, Figure 2; Hobson, 2009, p. 805).

In these ways, the various physiological states of sleep can be recognised and studied, but what about the experience? Approximately 14% of people report dreaming every night, 25% report dreaming frequently, and 6% never, and dream recall decreases with age (Blagrove, 2009). The emphasis here is on recall rather than dreaming itself because most dreams are never recalled. This was discovered in the 1950s when EEG studies first revealed the stages of sleep and people could be woken selectively at different stages.

When woken from non-REM sleep, people typically say either that nothing was going on in their mind or that they were thinking. As a simple example: 'I was asleep. I wasn't thinking about anything or dreaming about anything.' Or 'I was thinking about my nephew. It's his birthday soon and I must send him a card'. Non-REM reports are usually short and lacking in detail.

By contrast, when woken from REM sleep, people typically report that they were having complex, much longer, and often bizarre dreams; sometimes very bizarre, as in this excerpt:

I was at a conference and trying to get breakfast but the food and the people in line kept changing. My legs didn't work properly and I found it a great effort to hold my tray up. Then I realised why. My body was rotting away, and liquid was oozing from it. I thought I might be completely rotted before the end of the conference, but I thought I should still get some coffee if I had the strength.

We cannot say that this is a typical dream, for there is probably no such thing, but it has familiar elements that most people will probably recognise, especially the matter-of-fact response to profound bizarreness. The contents of dreams have been thoroughly studied using questionnaires and interviews, and by analysing reports using a scoring system originally developed in the 1960s by Calvin Hall and Robert van de Castle (1966; Domhoff, 1996). This counts such elements as settings, characters, emotions, social interactions, and misfortunes, all of which show remarkable consistency across times and cultures, with reliable sex differences, and reliable differences between the dreams of adults and children (Domhoff, 1996, Ch. 4).

For example, men dream more about other men than women do, and have more aggressive interactions. Children, by contrast, dream more often about animals, suffer more dream misfortunes, and are more often the victim of aggression than its initiator. Specific emotions or moods occur in about three quarters of dreams and are roughly equally positive and negative. Joy is the most common emotion reported, followed by anger and fear. Events in waking life often play a role in

'As for dreams – they're the "B-movies" of the mind – entertaining, but best forgotten'

(Horne, 2009, p. 709)

dreams, including dreams about such life events as surgery, psychotherapy, or marriage and divorce, as well as trivial events of the previous day.

If we wake up with a memory of dreaming, we are likely to try to make sense of what we dreamed; indeed, the sense-making process is part of the remembering. The natural tendency to attribute significance to dreams (perhaps even more than to events in waking life; Morewedge and Norton, 2009) was encouraged by Freud's (1900/1999) psychoanalytic approach to dream interpretation, which treats them as forms of wish fulfilment in which the real (or 'latent') content, deriving from the unconscious, is disguised in the superficial 'manifest' content of the dream scenarios. Jung (e.g. 1934–1936/1968) adapted these ideas to emphasise the role of basic archetypes that represent unconscious attitudes, and can be manifested in various dream symbols and figures which take dynamic forms depending on the dreamer and the dream context. Neither of these theories has stood the test of time. Although dream interpretation books and websites offering readymade templates for meaning-making are popular and many people believe their dreams give insight into unconscious beliefs and desires, there is no good reason to think that they do more than reflect current worries or hopes.

There are problems with generalising about dream content because of the effects of the method of collecting reports. For example, some researchers have asked people to keep dream diaries with dreams collected over long periods, while others ask just for the most recent dream. Selective reporting can be a problem with all collection methods, however, and the selection may take place at several stages: only some dreams are recalled on waking, some fade faster from memory after waking, and further selection can occur when people are asked to write a report or describe their dreams. In consequence, the occurrence of bizarre or interesting dreams may be exaggerated. Certainly many dreams are bizarre, but in studies that try to avoid selection problems, bizarreness is found in only about 10% of dreams.

This bizarreness takes different forms. Allan Hobson (1999) suggested three categories: *incongruity* involves the mismatching of features of characters, objects, actions, or settings; *discontinuity* involves sudden changes in these elements; *uncertainty* involves explicit vagueness. Research from his group suggested that the way characters and objects are transformed in dreams follows certain rules but that changes of scene and plot do not. Perhaps the strangest thing about

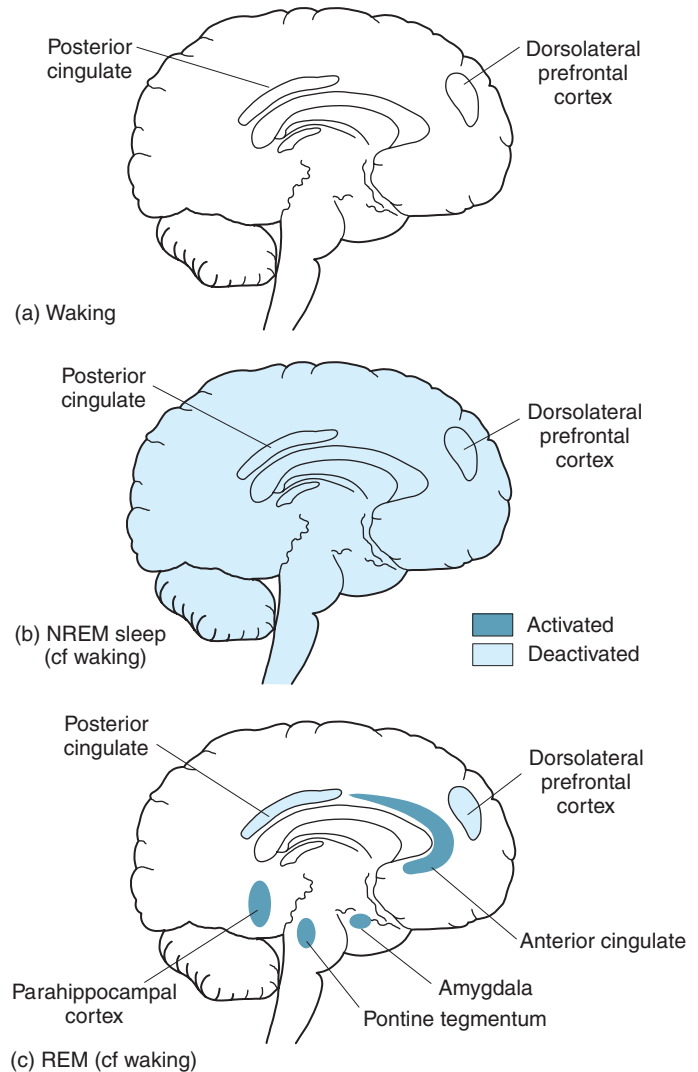


FIGURE 15.2 • Summary of PET study evidence of brain region activation in NREM and REM sleep. Compared with the blood flow distribution in waking (a), the global decreases observed in NREM sleep (b) suggest widespread deactivation consistent with the greatly diminished conscious experience early in the night. In REM sleep (c), many regions are activated about their levels in waking (dark blue), while others are deactivated (light blue; Hobson, 2002, p. 112).



PROFILE 15.1

Allan Hobson (b. 1933)



Known for his AIM model of dreaming states and his extensive work on sleep, Allan Hobson is both an experimental researcher and a psychiatrist, and is Professor Emeritus at Harvard Medical School. He began having lucid dreams after reading about them in 1962, and for decades he kept a dream journal. His dreams stopped after a stroke in 2001 but began again thirty-six days later, just as he began to walk again. He has long tried to understand the function of sleep, recently proposing that the brain optimises itself during sleep by minimising free energy and reducing the complexity of its model of the world. Hobson is a fervent critic of psychiatry's long reliance on psychoanalysis; he describes Freud's ideas as facile and erroneous, saying we have to wait for psychoanalysts to die since they will never recant. He has a dairy farm in Vermont where he has restored old buildings to house exhibitions and an art gallery. He is the author of many books on dreaming, including *The Dream Drugstore* (2001) and *Psychodynamic Neurology: Dreams, Consciousness, and Virtual Reality* (2015).

dreams is that while we are dreaming we rarely recognise how strange they are.

Finnish dream researcher Antti Revonsuo and his colleagues studied bizarre dreams in more detail using 592 dreams from the dream diaries of fifty-two students and measuring the bizarreness of their dream characters (Revonsuo and Tarkko, 2002). The most common type was bizarreness of dreamers' semantic knowledge about dream characters. Features intrinsic to the representation of a character were less often bizarre than the relationship between the character and the setting or the location, such as dreaming of 'the President having a cup of coffee in my kitchen'; there were also frequent changes, appearances, and disappearances of people and objects.

How can we make sense of all of this? Do these features of dreaming follow any patterns or rules? And if they do, can we understand them in terms of the underlying physiology?

FROM PHYSIOLOGY TO EXPERIENCE

Dream research seems to provide a perfect context in which to look for the neural correlates of consciousness. Various physiological, neurochemical, and behavioural variables can be correlated with subjective descriptions of dreams. On the surface, this might suggest the possibility of either reducing the experiences entirely to physical states, or equating the experiential with the physical, leading to the idea of just one combined objective/subjective space mapping and one concept of dreaming sleep, rather than two. This correlation between physiological states and subjective reports has supported decades of productive research into sleep and dreaming, and made it possible to map the three major states (waking, REM sleep, and non-REM sleep) in terms of their physiology. But does this help us to understand subjectivity or avoid the hard problem?

The best-known attempt at this sort of mapping is probably Hobson's AIM model (Chapter 13) depicting the idea of a unified 'brain-mind space'. The three states can be positioned in brain-mind space by measuring them along the three dimensions. Adding time as a fourth dimension, the values of A (activation energy), I (input source), and M (mode, or amine-choline ratio) all change and the process of cycling through the normal sleep stages can be represented by movement from one region of the space to another (Hobson, 2007). As in Tart's original conception for mapping altered states of consciousness (ASCs), large areas of the space remain unoccupied and the different states are discrete 'states of consciousness'. They



PROFILE 15.2

Antti Revonsuo (b. 1963)




As an undergraduate in psychology and philosophy, Antti Revonsuo wrote his thesis on how science-fiction stories present such traditional philosophical problems as soulless zombies and machine consciousness. From research on the cognitive deficits of neuropsychological patients, he moved into philosophy for a PhD on consciousness, combining philosophy, neural

are not points, however, but more like clouds in state space, especially the waking state in which the values are all high but change from moment to moment.

Things may not, however, be this simple. First, there is the obvious point that the map is crude, including only three dimensions, while the reality is much more complicated (LaBerge, 2000; Solms, 2000). This is not a serious problem since the scheme still provides a way of relating sleep states to other supposed ASCs and to the various neurotransmitter systems that control the overall state of the brain, and more detail and further dimensions could potentially be added.

More troublesome is that the correlation between REM and dreaming, while real enough, is not perfect. In the early days of sleep research, REM sleep and dreaming were often treated as equivalent, but subsequently people became more careful in referring either to the physical state or to the reported experiences. Dreaming is reported in about 70–95% of awakenings from REM sleep and roughly 5–10% of non-REM sleep, while mentation of some sort is reported from about 50% of non-REM awakenings. For example, after repetitive activities like playing a skiing game for two hours over a few days, many people woken shortly after falling asleep reported images clearly related to the game (Wamsley et al., 2010). In the skiing case, participants also reported images from past skiing experiences (especially crashes); if woken later in sleep, the imagery reported was more remote from the game, like stacking wood at a ski resort. The finding that non-REM experiences may become more ‘dreamlike’ as sleep continues raises the question of where exactly the boundary between ‘dreaming’ and ‘sleep mentation’ should be set: should we allow mentation to be any mental activity (e.g. perceptions, bodily feelings, thoughts) but try to restrict dreaming to ‘more elaborate, vivid, and story-like experiences recalled upon awakening’ (Kryger, Roth, and Dement, 2011, p. 585), or are such distinctions arbitrary and impossible to apply consistently? Maybe even REM and non-REM themselves cannot be neatly distinguished, and non-REM sleep might include covert REM processes (Nielsen, 2000). Overall, though, it is clear that being physiologically in REM sleep does not guarantee dreaming, and dreaming can occur without the physiological state of REM.

mechanisms, and ASCs. He has posts at both the University of Turku in Finland and the University of Skövde, Sweden. As a Harry Potter fan, he claims to use Professor Dumbledore’s Elderwand as a pointer, but we cannot confirm rumours that he was once a visiting professor at Hogwarts, specialising in Defence Against the Dark Arts. Revonsuo is best known for his evolutionary theory of dreaming as threat simulation, and his advocacy of the dreaming brain as a model for understanding consciousness. He describes himself as a ‘biological realist’, and believes that consciousness is a higher level of biological organisation in the brain. In future, he plans to work more on visual consciousness and anaesthesia.



CONCEPT 15.1 THE EVOLUTION OF DREAMING

- Why did dreaming evolve? The question here is not why *sleep* evolved. There are many competing theories about the evolutionary functions of REM and non-REM sleep in different species (Horne, 2006; Barrett and McNamara, 2012; Hobson and Friston, 2012), but the trickier question concerns dreams: do they have a function of their own or are they an inevitable concomitant of certain sleep states? As with the question of the evolution of consciousness itself (Chapter 11), we can find examples of all the main approaches.

Dreaming may have a crucial biological function. According to Antti Revonsuo’s Threat Simulation Theory, during most of human evolution serious physical and interpersonal threats meant a reproductive advantage for those who survived them, so dreaming evolved to simulate and practise dealing with these threats (an

argument also made for engagement with novels, drama, or films). Revonsuo (2000) shows that modern dreams include far more threatening events than people meet in waking life, and the dreamer usually engages appropriately with them. A broader view of 'dreaming as play' is proposed by Nicholas Humphrey (1983, 1986). Dreaming tests our physical, intellectual, and social skills and 'represents the most audacious and ingenious of nature's tricks for educating her psychologists' (1983, p. 85).

Flanagan argues that 'dreams are evolutionary epiphenomena' and have no adaptive function whatsoever. 'Dreaming came along as a free rider on a system designed to think and to sleep' (2000, pp. 100, 24). There is growing evidence that sleep plays an important role in reactivating and consolidating new memories, suggesting that the content and structure of dreams merely reflect these processes (Wamsley and Stickgold, 2011; Wamsley, 2014). The first theory to relate dreams to memory was Crick and Mitchison's (1983) proposal that neural networks become overloaded during learning and the function of REM sleep is to flood them to remove superfluous connections. In other words, we dream to forget. Hobson (2002) also connects dreaming to memory consolidation and considers dreams to be epiphenomenal, but on different grounds: that dream content has no significant influences on waking behaviour, and many people function perfectly well without recalling their dreams; while the REM state, by contrast, functions to minimise free energy and reduce the complexity of the brain's model of the external world (Hobson and Friston, 2012). Along similar lines, Tononi's 'synaptic homeostasis' hypothesis (Tononi and Cirelli, 2003) suggests that sleep regulates the excessive synaptic activation of wakefulness. Sleep is the price we pay for the brain's plasticity and dreams are a kind of play amongst its vast repertoire of memories.

The question of whether the conscious experience of dreaming plays a functional role or not is part of the wider question of whether consciousness in general has a function (Chapter 11), and remains unresolved. But either way, we can still use dreams in our own lives. Theories of dream interpretation, especially those based on Freud's psychoanalysis, have not stood the test of time (Webster, 1995; Hobson, 2002), but studying our own dreams can still reveal our motivations, hopes, and fears, encourage growing awareness, and even be a source of creativity and insight.

Further evidence comes from brain lesions. Damage to the ventral-mesial quadrant of the frontal lobe, which is involved in emotional motivation, or to the parietotemporo-occipital (TPO) junction, which is part of the sensory areas, reduces or obliterates dream recall while leaving REM sleep essentially normal (Solms, 2000). In other words, REM is neither necessary nor sufficient for dreaming.

Second, REM can occur when dreaming seems unlikely or even impossible. For example, human foetuses spend about 15 hours a day in REM sleep, babies spend less as they grow older, and children and adults less still. Yet foetuses cannot have anything like adult dreams because dreaming depends on prior experiences and on highly developed cognitive abilities which unborn babies lack. People with no visual experience, such as those born blind, dream without visual imagery but in words, ideas, and emotions, and in auditory, tactile, gustatory, and olfactory images. Dreams in people who become blind later in life gradually become less visual and more tactile (Meaidi et al., 2014). These people have plenty of experiences and a rich sense of self. But the new-born baby has neither.

As children grow older, their dreams closely reflect their developing cognitive abilities. Their dreams turn from rather static single dream images reported at age five or six, to more lively and dynamic imagery at age six or seven, with a dreamed self appearing only after the age of about seven years (Foulkes, 1993). We can therefore be sure that, whatever is going on for a foetus during REM sleep, it is not anything like an adult's dream. Hobson has speculated on this basis that what the mind-brain is doing in babies' REM sleep before dreaming appears is preparing itself for many integrative functions – and among these functions he includes consciousness. He suggests that REM sleep in early life before dreaming develops can be thought of as 'protoconscious' and as serving the purpose of allowing us to explore the possibilities and constraints of a virtual environment: 'The development of consciousness is thus seen as a gradual, time-consuming and lifelong process that builds on, and constantly uses, a more primitive innate virtual reality generator, the properties of which are defined for us in our dreams' (Hobson, 2009, p. 808). For Hobson, the ability to 'integrate' the dream state is what allows us to become aware of it.

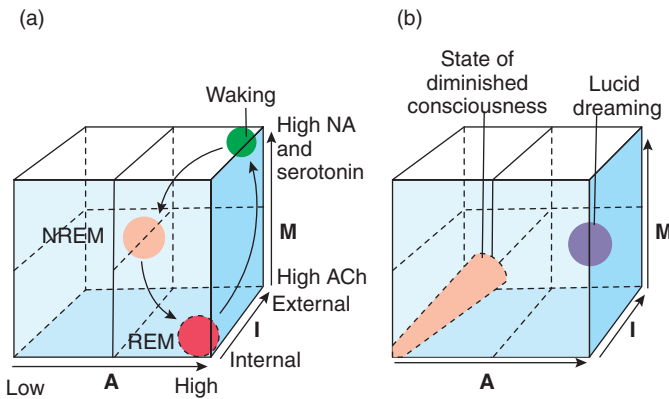


FIGURE 15.3 • AIM model of brain-mind state control. (a) The three dimensional AIM state-space model showing normal transitions within the AIM state space from waking to nonrapid eye movement (NREM) and then to rapid eye movement (REM) sleep. The x-axis represents A (for activation), the y-axis represents M (for modulation), and the z-axis represents I (for input-output gating). (b) Diseases, such as those that produce coma and minimally conscious states, occupy the left-hand segment of the space, owing to their low activation values. Lucid dreaming, which is a hybrid state with features of both waking and dreaming, is situated in the middle of the extreme right hand side of the AIM state space between waking and REM, towards either of which lucid dreamers are drawn (after Hobson, 2009, p. 808).

'Our conscious awareness during waking is an obvious adaptive advantage, but our conscious awareness during sleep may not be'

(Hobson, in Metzinger, 2009, p. 153)

Sleep in other species also seems likely to be very different from adult human sleep (Empson, 2001). Reptiles do not have REM sleep, but many birds and mammals do. Bottlenose dolphins, although extremely intelligent, do not seem to, and only one half of their brain sleeps at a time, in two-hour cycles, so they can keep watch for predators and know when to rise to the surface for air. REM-like sleep (with rapid eye movements, changes in body colouration, and arm-twitching) has been observed in cuttlefish though not octopuses (Frank et al., 2012). Mice and rats, dogs and cats, monkeys and apes all have REM sleep, and when we see their eyelids flickering or their whiskers twitching we can easily imagine that they are dreaming. But are we right to do so? We can guess, based on what we know of their cognitive abilities, that some of them might be enjoying complex visual and auditory images, perhaps even with narrative structure, but they cannot describe their dreams in words. So we cannot ask them and we cannot simply assume that REM equals dreaming.

Where can we go from here?

One possibility is that the physiology and the phenomenology can never be reduced to, or equated with, each other; that the fathomless abyss can never be crossed. Another possibility is that with further research, and better understanding of brain states and neurochemistry, we will learn exactly how brain states relate to the experience of dreaming.

There are already some hints in this direction. Dream contents have long been known to relate to eye movements, for example when someone reports having dreamt of watching a tennis match and distinct left-right eye movements are seen on the EEG recording. The same cortical areas appear to be involved in rapid eye movements as are involved in waking eye movements, and fMRI scans suggest that 'REMs are visually-guided saccades that reflexively explore dream imagery' (Hong et al., 2009).

The same sensory areas are activated when something is seen or heard as when it is imagined or remembered, and the same seems to be true of dreams. For example, the relative increases in activity in sensory areas and decreases in prefrontal areas are consistent with multisensory dreams lacking in executive control of action or decision-making. The emotions in dreams are consistent with increased activation in the amygdala, orbito-frontal cortex, and anterior cingulate, and the involvement of memory is related to activation of the hippocampus and connected areas (Maquet et al., 2005). Animal studies have revealed more about the connections between learning, memory, and dream content. For example, rats were trained to run on a circular track and activity in the hippocampus was recorded during the activity and when asleep (Louie and Wilson, 2001). Of more than forty REM episodes, about half repeated the unique signature of brain activity that was created as the animal ran. The correlation was so close that when the animal dreamed, researchers could reconstruct where it would be in the maze if it were awake and whether it was dreaming of running or standing still. More recently, studies of activity in the place cells of the hippocampus, which is precise enough to reconstruct a rat's position, have suggested that sleeping rats 'preplay' routes that they have seen will lead to food before actually exploring them, forming mental maps of the projected journey to and from the food (Ólafsdóttir et al., 2015).

Could we one day be able to deduce people's dreams from their brain activity? Scary as this prospect might seem, the first steps have already been taken. In the Gallant Lab at the University of California at Berkeley, scientists recorded many hours of fMRI data while people watched videos (Nishimoto et al., 2011) and created a huge 'dictionary' to relate the shapes, edges, and movements in the videos to activity at several thousand points in the viewer's brain. When they then showed a new video to the same person, they could use the dictionary to reconstruct a recognisable, if fuzzy, version of the video being watched. A similar method has since been applied to people sleeping inside a scanner and woken from REM sleep. By using the recorded data and the detailed dictionary, images of what they were dreaming about could be reconstructed (Horikawa et al., 2013). The computational power required was vast, but the principle has been proven: it should be possible to look at someone's brain activity and know what they are dreaming about. This is a huge step forward in our understanding, but perhaps only serves to make the gulf between physiology and experience seem more obvious.

Hobson's 'protoconsciousness' hypothesis about dreaming has recently been extended using ideas from predictive processing developed by theoretical neuroscientist Karl Friston. The function of sleep has long been hotly disputed, with theories ranging from maintaining neurotransmitter function to consolidating new memories, from driving metabolite clearance to promoting neural plasticity (Assefa et al., 2015). Hobson and Friston (2012) propose a new function. During sleep, the brain's 'virtual reality generator' (p. 85) simplifies its model of the waking world, so improving our ability to make reliable predictions. This idea is supported by various physiological observations. For example, pontine-geniculate-occipital (PGO) waves are involved in conveying eye-movement information within the visual system and might allow the brain to carry out predictive work during sleep. This could encompass both eye-movement command signals and the corollary discharge that allows us to predict the visual consequences of moving our eyes.

'REM sleep is a state of the brain that enables essential housekeeping functions, upon which waking consciousness depend[s]'

(Hobson and Friston, 2012, p. 87)

Unpredicted stimuli produce a robust spike (a startle response) in these waves in both waking and in REM sleep. Startle responses mean that a mismatch has occurred between what was predicted and what occurred, so their presence in REM sleep suggests that hypotheses are being tested against sensory evidence here, too. But a problem for this theory is that there is less habituation in this PGO startle response during sleep. This seems the wrong way round because surprise during dreams is very rare. Indeed, this is one of the oddities of dreaming: that we are so little surprised by truly bizarre dream events. The authors suggest that maybe we are incapable of surprise in dreams because what we experience is the result of top-down predictions in the thalamocortical system. This fits with the idea that dreaming is more similar to imagination (being driven by top-down intentions and predictions) than to sensory perception (driven in a bottom-up fashion by environmental percepts) (Nir and Tononi, 2010).

The brain, they say, is propelled in both sleeping and waking to infer the causes of its sensory sampling, like scientists driven to test their hypotheses (Hobson and Friston, 2014). The Cartesian theatre is a metaphor for the virtual reality models this creates. But there is no inner audience observing the show – only stories and fantasies being rehearsed and tested against sensory evidence. Is this a pernicious Cartesian theatre? Interestingly, Hobson and Friston arrive at a new kind of dualism, a duality between the conscious processes of inference and the physical brain states that encode them, claiming that this may help ‘dissolve some of the mysterious aspects of consciousness’ (p. 6).

Overall, however, Hobson and Friston are much more interested in the processing aspects of sleep than in the phenomenology of dreaming. For them, sleep is what does the important work, while dreams are merely ‘the subjective epiphenomena of the nocturnal products of our virtual reality generator and contain no new information’ (2012, p. 87). Maybe, they suggest, this is why it isn’t usually worth our while to remember them. But to dismiss something as important as dreaming as just an epiphenomenon raises its own philosophical problems (see Chapter 1).

Another example of how to connect physiology with experience comes from the dream bizarreness research by Revonsuo’s group. They argue that three types of bizarreness can be understood as failures of three types of binding: feature binding, contextual binding, and binding across time. They conclude that ‘more global forms of binding flounder much more frequently than those concerned with only local bundles of features’ and relate this to the number of distinct processing modules involved in generating different kinds of dream images (Revonsuo and Tarkko, 2002, p. 20). In other words, the harder it is for the brain to construct a certain kind of integrated image, the more likely it is that such an image will fall apart or show bizarre failures of binding during dreams.

This suggests that even the most peculiar of dream features may yield to a study of brain mechanisms during sleep. Even so, we are still relying on correlations, and as with all other aspects of conscious experience, we cannot say with confidence that dreaming and brain states are reducible to each other or are the same thing, nor can we confidently describe them in terms of ‘brain–mind states’.

So far we have been assuming that dreams are conscious experiences, but is this true? Some philosophers have questioned whether dreams are experiences at all (Malcolm, 1959; Dennett, 1976).

‘[Dreams are] the subjective epiphenomena of the nocturnal products of our virtual reality generator’

(Hobson and Friston, 2012, p. 87)

ARE DREAMS EXPERIENCES?

'dream consciousness is not normal consciousness, but it is consciousness nonetheless'

(Damasio, 2014, p. 111)

Of course dreams are experiences, you might say, and many would agree. The *Oxford English Dictionary* defines a dream as 'A series of images, thoughts, and emotions, often with a story-like quality, generated by mental activity during sleep' (January 2018). Psychology textbooks usually include dreams in sections on 'states of consciousness during sleep', and many philosophers and consciousness researchers accept this, too: 'Dreams are a form of consciousness, though of course quite different from full waking states' (Searle, 1997, p. 5); 'Dreaming is a subjective phenomenon of consciousness' (Revonsuo and Tarkko, 2002, p. 4); 'Dreams are conscious because they create the appearance of a world [...] Dreams are subjective states in that there is a phenomenal self' (Metzinger, 2009, p. 135); they are 'a second global state of consciousness aside from wakefulness' (Windt and Noreika, 2011), or 'an altered state of consciousness that is difficult to recall in waking' (Hobson, 2014, p. 4). Hobson defends the commonsense view like this: 'Our dreams are not mysterious phenomena, they are conscious events. Here's the simplest test: Are we aware of what happens in our dreams? Of course. Therefore, dreaming is a conscious experience' (Hobson, 1999, p. 209).

But are we really aware in our dreams? Suppose that I wake from a dream and think, 'Wow, that was a weird dream. I remember I was trying to get some coffee'. At the time of waking, I seem to *have been having* the dream. Indeed, I am completely convinced that a moment ago I was dreaming of being in the cafeteria, even if the details slip quickly away and I cannot hang onto them, let alone report them all. But there are some serious problems here.

Some concern the self. Although I am sure that 'I' was dreaming, the self in the dream was not like my normal waking self. This strange dream-self didn't realise she was dreaming; she accepted that the people and the food kept changing in impossible ways, showed little disgust or surprise at the state of her body, and in general treated everything as though it was real. Was it really me who dreamt it? Maybe not – but perhaps, as Metzinger would argue, this does not matter because there was *some kind of* phenomenal self in the dream and this is enough to support a PSM, a phenomenal self-model.

Other problems concern the lack of insight during dreams. Taking Tart's subjective definition of an ASC (Chapter 13), there is clearly 'a qualitative alteration in the overall pattern of mental functioning', but unlike in most drug-induced states, or during sensory deprivation or starvation, it is not true that 'the experiencer feels his consciousness is radically different from the way it functions ordinarily'; the experiencer, at least in non-lucid dreams, fails to notice this 'radical' change. So, by this definition, we are forced to the curious conclusion that the ordinary dream, that most classic of all ASCs, is not really an ASC at all. Oddly enough, by the same definition, a lucid dream *is* an ASC because now the experiencer *does* realise it is a dream.

Other peculiarities concern the status of the dream: if I start to doubt whether I really did have that dream, the only evidence to call on is my own memories, and those are vague and fade fast. One response to such doubts goes back to 1861, when French physician Alfred Maury described a long and complicated dream about the French revolution, culminating in his being led to the guillotine.

Just as his head came off, he awoke to find that the headboard had fallen on his neck (Maury, 1861, pp. 133–134). He proposed that dreams do not happen in real time but are entirely concocted in the moment of waking up. This theory became popular, perhaps because so many people have the experience of dreaming about a church bell ringing or a wolf howling, only to wake to the sound of their alarm clock. It is also psychologically plausible in the sense that humans are very good at constructing stories and quick at confabulating. But it is not true.

In the 1950s, people sleeping in the lab were asked to describe their dreams and they gave longer descriptions the longer they had been in REM sleep. Other experiments tried incorporating external stimuli into dreams. Sounds, taps on the skin, flashes of light, and drips and sprays of water have all been used, and when they don't wake the sleeper they can sometimes influence dream content, allowing dream events to be timed. These results, as well as the animal studies described above, show that dreams take about the same time as would waking events. All this suggests that dreams are not concocted in a flash on waking up, but really do take time.

Other responses to these doubts are more subtle. Dennett provides a selection of fanciful theories playing with the relationship between experience and memory. On the 'cassette theory of dreams', the brain holds a store of potential dreams recorded and ready for use. On waking from REM sleep, a 'cassette' is pulled out of storage, to match the sound of the alarm clock if necessary, and hey presto, we seem to have been dreaming. On this theory, there are no real dreams. There are no events or images presented 'in consciousness', but only recollections of dreams that were never actually experienced. 'On the cassette theory it is not like anything to dream, although it is like something *to have dreamed*'. On the cassette theory, dreams are not experiences we have during sleep' (1976, p. 138).

The point of this theory is not that it might be literally true (even if we update the cassette to an MP4), but that it provides the basis for another possibility – that the equivalent of cassette dreams might be composed during the REM period prior to waking. We can now compare the normal theory that dreams are conscious experiences during sleep, and the new theory that dreams are composed unconsciously during sleep and then 'remembered' on waking up. The question is this. Could we ever tell which was right?

The answer seems to be no. It is no good asking dreamers whether their dreams really occurred 'in consciousness', because all they have is their memories and they will always say 'yes'. And it is no good looking inside their brain because even if we could see the neural events that correlate with imagining cups of coffee or trying to walk, we still have no way of finding out whether those neural events were 'in consciousness' or not. There is no special place in the brain where consciousness happens, or, in terms of Dennett's later theory (1991), there is no Cartesian theatre in which the dreams either were, or were not, displayed. We are left with two theories that seem empirically indistinguishable, so this is again 'a difference that makes no difference.'

But the tendency to distinguish conscious from unconscious elements of dreaming remains common – even amongst researchers who claim consistency with multiple drafts theory. One account models the phenomenology of dreaming as a feedback system involving memories, interpretive processes brought to bear on them, and

'The conscious output of the dream is what will be recalled by the dreamer'

(Cicogna and Bosinelli, 2001, p. 38)




FIGURE 15.4 • According to the retro-selection theory (Blackmore, 2004), dreams are not conscious experiences. They are concocted retrospectively on waking by selecting from the myriad trains of thoughts and images that were going on in parallel in the dreaming brain. So on waking, this dreamer might recall that he had ripped some flowers from their pot, rushed off on skis to escape retribution, arrived in a forest and had a picnic and a bottle of wine under a pine tree. With many more parallel processes going on than are shown here, a very large number of potential dreams are possible, and alarm clocks ringing or other sounds on waking might easily influence which was selected.

monitoring of phenomenal experience, which together plan and co-create the dream: 'As it develops, the unconscious planning simultaneously becomes the conscious syntactic organization of the dream' (Cicogna and Bosinelli, 2001, p. 34).

'it is not like anything to dream, although it is like something to have dreamed'

(Dennett, 1976, p. 138)

The idea is that 'the iterative feedback mechanism constructs successive drafts of the dream [. . .] at an unconscious level with only the end product, the dream itself, being accessible to awareness' (p. 34). The top-down processes make the first draft, the memory elements activate or inhibit other elements, and so different versions are created. Although the authors agree with Dennett in rejecting the idea of a central controller, they still describe the operations involved in dream generation as 'unconscious', and so end up having to ask what the function is of the conscious processes in the dream. And because only the final product is 'conscious', we must also still ask what it is that makes the difference.

Returning to the question of timing, there is still an interesting conflict in the findings described above. On the one hand, we know that dreams occur in real time; on the other, we know that people often wake from dreams in which the event that woke them fits the end of a long dream story. How can this be?

One way of explaining this, very much in the spirit of multiple drafts theory, is the retro-selection theory of dreams (Blackmore, 2004). During REM sleep numerous brain processes are going on at once, none of which is either in or out of consciousness. On waking up, a story is concocted by selecting one out of a vast number of possible threads running through the multiple and confusing scraps of memory that remain. The chosen story is woven backwards to fit the timing, but is only one of many such stories that might have been selected had a different event woken the dreamer up. The important point is that there is no version of the story that counts as the actual dream, consciously experienced at the time. This theory can resolve the peculiar conflict described above, but it means accepting that there is no right answer to the question 'what did I really dream about?' The theory is testable. For example, since lots of brain events are going on during sleep, these should be observable using the methods developed by Tomoyasu Horikawa and others (2013). And it ought to be possible to wake dreamers using different types of stimulus and

expect them to report different dream contents to match because what counts as the dream story was determined only on waking. On this theory, dreams are not conscious experiences until we have woken up and decided they were. They do not happen 'in consciousness', but then nor does anything else (Blackmore, 2014).

This theory gives us an alternative to both the standard theory that dreams are conscious experiences happening during sleep, and the alternative that dreams are composed unconsciously during sleep and then 'become conscious' on waking up, both of which have serious theoretical and empirical drawbacks.

'When the sleeper awakes, a dream is concocted, backwards, by selecting any one of the possible multiple threads'

(Blackmore, 2004)

THE BORDERS OF SLEEP

Strange, dreamlike experiences can happen before we fall asleep or as we are waking up. At these times, when sensory input is reduced, hallucinations are common, and range from simple visual forms or musical notes to sensations on the skin or imagined changes in the location of a limb. This type of hallucination was first described in 1848 by Maury, who called those which happen while falling asleep hypnagogic images or hypnagogic hallucinations and those that occur on waking hypnopompic images (Mavromatis, 1987).



PRACTICE 15.1 STAYING AWAKE WHILE FALLING ASLEEP

It is easy to start to explore the borderland between reality and imagination by learning to hover on the edge of sleep. Do this exercise for a week and you may be rewarded with fascinating hallucinations and insights. The visions and sounds may be frightening for some people, and you should not pursue this if you find it too unpleasant.

Go to bed as usual, lie in your normal position, but then try to keep your mind clear and empty. When any thoughts arise, gently let them go, as you did when practising meditation. Look into the darkness in front of you and watch for patterns. Listen attentively for sounds. When you see or hear things, or feel odd twitches in your muscles, try not to react but stay relaxed and keep watching and listening.

There are two difficulties. The exercise may keep you awake when you want to sleep, or force you to have a clear mind when you would rather indulge in fantasy or worry. We can only suggest that the visions may be worth the loss of sleep, and that in fact you will not take much longer to go to sleep than normal, however it feels.

Alternatively, you may find that you drop off to sleep too fast. One suggestion from the Western occult tradition is to lie on your back, holding one forearm vertical. As you fall asleep, the arm drops and wakes you. This way you can oscillate between sleep and waking. In any case, lying on your back makes hypnagogic imagery and sleep paralysis more likely. Like many of these exercises, this one gets rapidly easier with practice.

As with other kinds of hallucination (Chapter 14), the form constants are common, and many people describe flying or falling through tunnels, tubes, or cones, or through black spaces lit by stars. They see whirling circles or suns, luminous points or streaks, and vibrating coloured threads. More rarely people see animals, people, mythical creatures, or complex landscapes, or they hear chattering and muttering voices. Sometimes people who have been doing something for many hours in the day see perseverative images of those things as they fall asleep, such as weeds if they have been gardening, or endless shoals of fish if they have been snorkelling. Others hear their own name being called distinctly as they fall asleep, and this can be so realistic that they get up to see who is there. A few people learn to control their hypnagogic images, but they say it is more like ‘wishing’ than ‘willing’ because you don’t always get what you wanted (Mavromatis, 1987). Mostly the experiences are vivid and uncontrollable, and are not mistaken for reality as the strictest definition of hallucination requires.

These hallucinations can be combined with one of the oddest phenomena on the borders of sleep: sleep paralysis (SP), illustrated in the example at the end of Chapter 14.

9 December. I had a dream from which I awoke with a throbbing heart. I saw as if I were in Moscow in my house, in the big sitting room, and Joseph Alexéevich came in from the drawing room. It was as if I knew at once that the process of regeneration had already taken place in him, and I rushed to meet him. It was as if I embraced him and kissed his hands, and he said, ‘Have you noticed that my face is different?’ I looked at him, still holding him in my arms, and it was as if I saw that his face was young, but that he had no hair on his head and his features were very different, and as if I said, ‘I should have known you had I met you by chance’, and thought to myself, ‘Am I telling the truth?’ And suddenly I saw him lying like a dead corpse; then he gradually recovered and went with me into my study carrying a large book, decorated with Alexandrian senna. It was as if I said, ‘I drew that’, and he answered by bowing his head. I opened the book, and on all the pages there were excellent drawings. It was as if I knew that these drawings represented the love adventures of the soul with her beloved. [...] As if looking at those drawings I felt that I was doing wrong, but could not tear myself away from them.

(Leo Tolstoy, *War and Peace* [Война и мир], Book VI, Ch 10, our translation, 1869)

STRANGE DREAMS

When Sue realised that the ski lift, mountains, and rising sun were all a dream, she knew she could fly, and soared up into the cold morning air over the mountain peaks. Flying dreams are reported



ACTIVITY 15.1

Discussing hypnagogia

The exercise in Practice 15.1 lends itself well to group work. Ask everyone to practise ‘staying awake while falling asleep’ for several days, to keep a pencil and paper by the bed, and to write down anything they experience. It may be impossible to record the experiences immediately when they happen because the most interesting ones happen right on the edge of sleep, but they can be written down, or drawn, in the morning. Ask participants to bring any notes and drawings to the discussion.

Were there common themes? Are the form constants discernible in the descriptions? Is there any pattern to who did and did not have hallucinations? Did anyone experience sleep paralysis or body distortions? Was the experience pleasurable?

by about half the population, and are usually pleasant or even joyful. Falling dreams are also common and sometimes end with a myoclonic jerk – an involuntary muscle spasm that occurs during the shift from waking into sleep. Most such dreams are not lucid – that is, the dreamer rarely thinks ‘Wow, I can’t fly in normal life, so this must be a dream’ – but they can sometimes alert people to their state and lead to lucidity.

Another odd dream is the ‘false awakening’, a dream of having woken up. Sometimes everything looks quite normal and so the dreamer gets on with dressing and eating breakfast until he really wakes up and has to start all over again. A famous example was described by the French biologist Yves Delage in 1919. Delage was asleep when he heard a knock at the door. He got up to find a visitor asking him to come quickly and attend to a sick friend. He leapt up, dressed, and started to wash, whereupon the cold water on his face woke him up and he realised it was only a dream. Back in bed, he heard the same voice again and, fearing he must have fallen asleep, leapt out of bed and repeated the dressing and washing four times before he really woke up (Green, 1968a).

In other false awakenings, people report greenish light, glowing objects, eerie feelings, and humming or buzzing sounds. These are all reminiscent of hypnagogic experiences, and prompt the odd thought that it may sometimes be impossible to know whether one is awake and hallucinating, or only dreaming one is awake. In the first case, the bedroom is real even if the hallucinations are not, but in the second, the whole room and everything in it is dreamed. Experiences like this, in which the whole environment is replaced by hallucinations, are sometimes called ‘metachoric experiences’ (Green and McCreery, 1975). This profound doubt can extend over crisis apparitions, fairy abductions, alien visitations, and even some drug experiences. Without physiological monitoring we cannot know whether the person had their eyes open, as often claimed, or was fast asleep (Blackmore, 2017).

Then a terrible thought occurred to her. What if this was still a dream? What if she had only dreamed that she had really woken up? How could she tell? She pinched herself, hard. She felt the pinch all right and saw she’d made a bright red mark on her skin but then she



CONCEPT 15.2

SLEEP PARALYSIS

The experience described at the end of Chapter 14 is a typical account of sleep paralysis (SP), derived from hundreds of cases gathered via magazine advertisements (Parker and Blackmore, 2002). SP is one symptom of the serious sleep disorder narcolepsy, and for that reason may be treated as pathological, but SP is common in healthy people, with estimates coming from Canada (21%), Hong Kong (37%), Japan (40%), Nigeria (44%), England (46%), and Newfoundland (62%) (Parker and Blackmore, 2002). One overview of thirty-five previous studies estimated a lifetime incidence of 8% in the general population, 28% among students, and 32% of psychiatric patients (Sharpless and Barber, 2011).

SP most commonly occurs during sleep onset REM (SOREM) and can be thought of as an intrusion of REM into either light sleep or waking (Nelson et al., 2006). The person feels awake but the voluntary muscles are paralysed. The most common features are fear, the ‘sense of presence’ (often evil or frightening), humming, buzzing, or grinding noises, pressure on the chest, vibrations through the body, touches on the limbs, and sensations of floating or even out-of-body experiences (Cheyne, Newby-Clark, and Rueffer, 1999; Blackmore, 2017; Denis and Poerio, 2017). Many people are terrified because they believe that the presence is a real ghost or alien, or because they think they must be going mad. Knowing something about SP makes it much less frightening.

SP can, with difficulty, be induced in the laboratory by repeatedly waking people just after they have entered REM, keeping them awake for an hour, and then letting them sleep again (Inugami and Ma, 2002). Most features of SP have been independently induced by transcranial magnetic stimulation, in particular by stimulation of the temporal lobes (Persinger, 1999). For example, the sense of presence is thought to be a displaced version of one’s own body schema and can be induced by stimulation of the left temporoparietal junction (Arzy et al., 2006; Brugger, 2006).

Some regular experiencers learn to prevent SP by avoiding sleeping on their back and getting regular sleep. When it occurs, the best way to cope is just to relax and wait for it to stop, which it usually does within a few seconds, although it is difficult to follow this advice if you are terrified. Other methods include trying to move just a little finger or toe, or blinking rapidly.

Many cultures have sleep paralysis myths, such as the incubus and succubus of medieval lore, and the seductive Babylonian Lilitu or demoness of the wind. The 'Old Hag' of Newfoundland is 'The terror that comes in the night' (Hufford, 1982), sitting on victims' chests and trying to suffocate them. The same experience is called Kanashibari (meaning 'to tie with an iron rope') in Japan, Ha-wi-nulita (or being squeezed by scissors) in Korea, and Kokma (attacks by the spirits of unbaptised babies) in St Lucia. The latest SP myth may be alien abductions, which include all the usual features of paralysis, suffocation, floating sensations, sense of presence, touches on the body, and vibrating or humming noises. It seems that peoples in many times and places have invented myths and entities to account for this common physiological occurrence. For a personal account, read Ronald Siegel's chapter 'The succubus' (1992, pp. 83–90).

realised she might just have dreamed the feeling and the mark. So that was no proof. She banged her hand hard on the bedside table. It felt solid enough and the lamp jumped and wobbled and nearly fell over but the dream might have invented the table and the lamp and made them convincing enough to seem real. How could she tell?

I know, thought Jinny. I know what to do. In dreams you can fly. And she remembered all the times she'd flown in her dreams; flying with Hatty in the great blue sky, flying over the sea and above the boats, flying through forests without being seen. This will prove it, thought Jinny. I'll see if I can fly. So she climbed up on the bed and flapped her arms. Nothing happened. She jumped up and down and nothing happened. She lay on her tummy and swam with her arms and nothing happened. I think I really am awake this time, she thought. But she fell asleep still wondering. Could she have dreamed that she couldn't fly? Could she have dreamed that she was dreaming that she was trying to find out if she was dreaming? Could she?

(Sue Blackmore, *Jinny Jana's Giant Journeys*, 2016)

Why don't we realise we are dreaming at the time? This is the oddest and most frustrating thing about ordinary dreams: that we can fly, drive a Porsche across the sea, or survive the devastation of an atom bomb, with no insight at all. Sometimes, however, critical doubt does creep in, prompted by strong emotions, by incongruities in the dream, or by recognising recurring themes from previous dreams (Green, 1968a; Gackenbach and LaBerge, 1988). If we ask the question 'Am I dreaming?', we are having what the English psychologist and pioneer of lucid-dream research Celia Green calls a 'prelucid dream'. Even then, it is common for dreamers to give the wrong answer. There are accounts of people asking dream characters whether they are dreaming, splitting into two and arguing over whether they are dreaming, or trying to pinch themselves to find out. Of course, the pinching test fails for those who dream a dream pinch and feel a realistic dream pain.

LUCID DREAMS

When the correct conclusion is reached, the dream becomes a lucid dream, 'a global simulation of a world in which we suddenly become aware that it is indeed just a simulation' (Metzinger, 2009, p. 140) – a tunnel whose inhabitant realises it is a tunnel.

'In lucid dreams, part of the brain operates in the primary mode while another has access to secondary consciousness'

(Voss et al., 2013, p. 9)

This realisation can have extraordinary consequences. Not only do people describe lucidity as like ‘waking up in the dream’ or ‘becoming conscious while dreaming’, but many claim that once lucid they can fly or float, take charge of the course of their dream, or change the objects and scenery at will. ‘The subject of a lucid dream is not a passive victim lost in a sequence of bizarre episodes but rather is a full-blown agent, capable of selecting from a variety of possible actions’ (Metzinger, 2009, p. 143). As Hobson puts it, ‘part of my brain-mind wakes up and [. . .] then I can have a lot of fun. I can watch the dreams [. . .], I can influence the dream content’ (Hobson, 2002, p. 142).

The shift from ordinary to lucid dreaming has been characterised in many ways. For Hobson, Voss, and colleagues, there is a difference between ‘primary’ and ‘secondary’ consciousness. Primary consciousness is what we have in normal dreams. It is governed by what is immediately present; all we can do is cope with the immediate and constantly changing scenery, rather than influencing the ongoing experience. When we wake up, we enter higher-order consciousness, in which we can plan ahead, reflect on the past, and contemplate the future. When we ‘wake up’ into a lucid dream, they say, ‘part of the brain operates in the primary mode while another has access to secondary consciousness’ (Voss et al., 2013, p. 9). But what does it mean for one part of a brain (or brain–mind) to operate in a different mode from another, and to have access to one kind of consciousness or not?

Another way of describing this shift is to say that in ordinary dreams we have ‘phenomenal awareness’ (of objects and events) and ‘self-awareness’, but that only in lucid dreams do we also have ‘meta-awareness’ (awareness of one’s own mental activity) (Cicogna and Bosinelli, 2001). Lucid dreams may feel like they involve this sort of change, but do these distinctions really work? It is all too easy to create hierarchies of states and levels of consciousness, but they rarely stand up to careful scrutiny. Retro-selection theory requires no such distinctions. It simply implies that in a lucid dream, instead of waking up and only then constructing a story out of the remaining threads of memory, the selecting and story-constructing is done during the dream.

Note that although it feels as though the increased consciousness causes the ability to control the dream, this conclusion is not warranted by the correlation. All we know is that in lucid dreams critical thinking, dream control, flying, and the sense of being more awake, or more conscious, or more ‘myself’, all occur together. We also know that in lucid dreams people report more insight, logical thought, control over thoughts and actions, and positive emotion than in non-lucid dreams (Voss et al., 2013). But there doesn’t seem to be a strong correlation between ‘thought’ and ‘insight’ – that is, knowing you are dreaming isn’t necessarily related to thinking logically about it or finding it more realistic.

The term ‘lucid dream’ was coined by the Dutch psychiatrist Frederik van Eeden in 1913, and although the name does not describe this kind of dream at all well



FIGURE 15.5 • David Howard suffered from narcolepsy, a sleep disorder characterised by periods of sleepiness or sudden sleep during the day, as well as abnormalities of dreaming sleep and hallucinations. During narcoleptic episodes he claimed to have been frequently abducted by aliens, operated on by them, and taken to their ships and planets. His paintings show the rich details of his memory for these experiences.

‘Lucidity involves the cognitive realization that you are currently dreaming [. . .], not necessarily experiencing your dreams as unreal or as a merely virtual reality’

(Voss et al., 2013, p. 19)



ACTIVITY 15.2

Inducing lucid dreams

As a class activity, divide the group into three and give everyone a week to try to have a lucid dream. It is best to assign people randomly to the groups, but if you have several good lucid dreamers in the class, spread them equally across the groups. Compare the number of lucid dreams achieved in each group and discuss the results. (If you have enough data, use ANOVA based on the number of lucid dreams per participant. Alternatively, compare two groups using an independent t-test.) Even if the groups are too small for statistical analysis, the experiences of trying, the frustrations of failing, and the pleasures of successful lucidity will provide plenty of scope for discussion.

The groups are as follows:

- 1 **Control group.** Use no special technique. People often report having lucid dreams after simply hearing or reading about them, so this group provides a better baseline than people's previous levels of lucidity. If you have fewer than about thirty participants, drop this group and use only 2 and 3.
- 2 **Daytime awareness.** Use letters drawn on the hands as in Practice 15.2.
- 3 **Nighttime intention.** The idea is to go to sleep thinking about dreams and intending to notice the next time you have one. Before you fall asleep at night, try to remember the dream you had the night before, or any recent dream. Go through your memory noticing odd features, the way things behaved, or anything that is characteristic of your dreams. Tell yourself, 'Next time I dream this, I will realise I'm dreaming'.

A more arduous version of this is LaBerge's MILD (mnemonic induction of lucid dreaming) technique (for more details see LaBerge, 1985; LaBerge and Rheingold, 1990). Wake yourself with an alarm in the early hours of the morning. If you have been dreaming, mentally rehearse the dream or, better still, get up and write it down. As you go to sleep again, visualise yourself back in the dream, but this time you realise it is a dream. Keep rehearsing the dream until you fall asleep.

('lucid' means either clearly expressed, or bright or luminous), it has stuck. Surveys show that about 50% of people claim to have had at least one lucid dream in their lives, and about 20% have one a month or more. This figure may be unreliable, because although lucid dreamers will easily recognise the description, people who have never had a lucid dream may misunderstand it. With this proviso, surveys show no correlations with age, sex, personality measures, or basic demographic variables, but the same people tend to report lucid dreaming, flying and falling dreams, and out-of-body experiences (Green, 1968a; Blackmore, 1982; Gackenbach and LaBerge, 1988).

Lucid dreams were long considered beyond the pale of serious sleep research, and were studied only by psychical researchers and parapsychologists. Even in the mid-twentieth century, many psychologists rejected the whole idea, arguing that self-reflection and conscious choice are impossible in dreams, so lucid dreams must really occur before or after sleep, or during micro-awakenings.

They were proved wrong. The breakthrough was made simultaneously and independently by two young psychologists, Keith Hearne at the University of Hull in England, and Stephen LaBerge at Stanford University in California. The problem they faced was simple. In REM sleep the voluntary muscles are paralysed, so a dreamer who becomes lucid cannot shout out 'Hey, listen to me, I'm dreaming' or even press a button to indicate lucidity. What Hearne and LaBerge realised was that dreamers could still move their eyes. In Hearne's laboratory, Alan Worsley was the first oneironaut (or dream explorer) to signal from a lucid dream. He decided in advance to move his eyes left and right eight times in succession whenever he became lucid and Hearne picked up the signals on a polygraph. He found them in the midst of REM sleep (Hearne, 1978), a finding that has been confirmed many times since (LaBerge, 1990).

Further research has shown that lucid dreams last an average of two minutes, although they can last as long as fifty minutes. They usually occur in the early hours of the morning, nearly half an hour into a REM period and towards the end of a burst of rapid eye movements. Onset tends to coincide with times of particularly high arousal during REM sleep and is associated with pauses in breathing, brief changes in heart rate, and skin response changes. There is also increased activity in the left parietal lobe, which may be related to the more solid sense of self in lucid dreams (Holzinger, LaBerge, and Levitan, 2006) and more 40 Hz power than in ordinary dreaming, especially in frontal regions (Voss et al., 2009). Stimulating REM sleepers' brains with 40 Hz currents has also been shown to induce lucid dreaming, especially in those who have had lucid dreams before, and with strong correlations



PRACTICE 15.2

BECOMING LUCID

If you are taking part in the class activity (Activity 2) try whichever induction technique is assigned to you. Otherwise practise this one.

Take a pen and write a large **D** on one hand, for Dreaming, and a large **A** on the other, for Awake. As many times as you can, every day, look at these two letters and ask '**Am I awake or am I dreaming?**' If you get thoroughly into the habit of doing this during the day, the habit should carry over into sleep. You may then find yourself looking at your hands in a dream and asking 'Am I awake or am I dreaming?' This is a pre-lucid dream. All you have to do is answer correctly and you're lucid.

Did it work? What happened in the dream? What happened to your awareness during the day?

between levels of 40 Hz activity and ratings of insight (awareness of being in a dream) and dissociation (experiencing the dream from a third-person perspective) (Voss et al., 2014). Given claims relating 40 Hz power to consciousness, this also fits with the notion that lucid dreaming is a state hovering between waking and dreaming sleep.



FIGURE 15.6 • How can you test whether you are dreaming? In 1920s London, Oliver Fox made many such tests during his experiences of astral projection and lucid dreaming. 'I dreamed that my wife and I awoke, got up, and dressed. On pulling up the blind, we made the amazing discovery that the row of houses opposite had vanished and in their place were bare fields. I said to my wife, "This means I am dreaming, though everything seems so real and I feel perfectly awake. Those houses could not disappear in the night, and look at all that grass!" But though my wife was greatly puzzled, I could not convince her it was a dream. "Well", I continued, "I am prepared to stand by my reason and put it to the test. I will jump out of the window, and I shall take no harm." Ruthlessly ignoring her pleading and objecting, I opened the window and climbed out on to the sill. I then jumped, and floated gently down into the street. When my feet touched the pavement, I awoke. My wife had no memory of dreaming.' (Fox, 1962, p. 69).




FIGURE 15.7 • Signal-verified lucid dream. Four channels of physiological data (central EEG [C3-A2], left and right eye movements [LOC and ROC], and chin muscle tone [EMG]) from the last 8 min of a 30-min REM period are shown. On awakening the sleeper reported having made five eye-movement signals (labelled 1-5 in figure). The first signal (1, LRLR) marked the onset of lucidity (LaBerge, 2000, Fig. 1).

The signalling method means we no longer have to rely on retrospective verbal report, and so allows us to answer some classic questions about dreams. Correlations between dream content and physiology can now be timed accurately, and lucid dreamers can be given pre-sleep instructions to carry out particular activities during their dream and signal as they do so. One example is that question about how long dreams last. Lucid dreamers can accurately estimate the time taken by dreamed events, and when asked to count to ten during lucid dreams and again during waking, they took about the same length of time (LaBerge, 2000). But physical actions take longer. In one study, dreaming of doing squats took 40% longer than physically doing them (Erlacher and Schredl, 2004). In another, lucid dreamers walked 10, 20, or 30 steps and did a short gymnastics routine; both these actions took longer in the dream than they would in real life (Erlacher et al., 2014). Respiration and heart rate rose when doing squats in a lucid dream (Erlacher and Schredl, 2008) and when performing different actions, the muscles that would be used for those actions in waking life twitched slightly during the dreams. Pre-agreed voluntary breathing patterns coincide with actual breathing, and in one study a woman's erotic lucid dream coincided with actual sexual arousal and a measurable orgasm (for a review see LaBerge, 1990).

Could practising a skill during a lucid dream improve that skill in waking life? In a survey of hundreds of German athletes, over half reported having lucid dreams and nearly 10% claimed to use lucid dreaming to practise their sport (Erlacher, Stumbrys, and Schredl, 2012). In experiments testing simple skills such as

finger-tapping or throwing coins into a cup, practice worked better in lucid dreams than when awake (Stumbrys and Erlacher, 2016).

Another question is whether the eye movements of REM sleep correspond to dream events. This had been suspected from observations of non-lucid dreams but is easily confirmed with expert lucid dreamers who can deliberately do things like playing tennis, revealing that eye movements do indeed reflect dream events. Furthermore, experiments tracking moving objects during the dream revealed that lucid-dream eye movements more closely resemble the smooth pursuit of waking vision than the saccadic eye movements associated with imagination (LaBerge, 1985, 1990). But smooth pursuit is now known to occur in mental imagery, too, especially with increased drowsiness (de'Sperati and Santandrea, 2005), suggesting a fluid perceptual continuum between imagining, dreaming, and lucid dreaming.

Few people can induce lucid dreams at will, but there are techniques that can help. Several machines work on the twin principles of first detecting REM sleep and then delivering a stimulus strong enough to increase arousal slightly, but not strong enough to wake the sleeper, including Hearne's (1990) Dream Machine and LaBerge's DreamLight. Of forty-four participants who used the DreamLight in the laboratory, 55% had at least one lucid dream and two had their first ever lucid dream this way (LaBerge, 1985). The later NovaDreamer packed all the hardware into goggles which could be worn at home. Competitors including the REM Dreamer have added features like interactive controls, and lucid-dreaming apps now claim to detect dream sleep via mattress movement, so all you need is your phone. Apps like psychologist Richard Wiseman's Dream:ON also offer a range of soundscapes to help shape your dream into a peaceful garden or ocean scene. It wakes you up once it detects your dream is over, and asks you to submit a dream report to its 'dream catcher' database.

Other methods include maintaining awareness while falling asleep, LaBerge's MILD technique (see Activity 2), and other procedures that increase awareness during the day rather than just at night. These are based on the idea that we spend much of our time in a waking daze and if we could only be more lucid in waking life, it might carry over into dreaming. These methods are similar to the age-old techniques of meditation and mindfulness. Indeed, advanced practitioners of meditation claim to maintain awareness through a large proportion of their sleep, and research has found associations between practising meditation and increased lucidity (Gackenbach and Bosveld, 1989). Some people even choose to meditate once they become lucid, echoing the ancient Tibetan practice of deepening meditative insight through lucid 'dream yoga'. Others choose wish fulfilment, problem-solving, skill-training, and mental or physical healing (Stumbrys and Erlacher, 2016).




FIGURE 15.8 • Dream time estimations. LaBerge asked participants to estimate ten-second intervals by counting, 'one thousand and one, one thousand and two, etc.' during their lucid dreams. Signals marking the beginning and end of the subjective intervals allowed comparison with objective time. In all cases, time estimates during the lucid dreams were very close to the actual time between signals (LaBerge, 2000, Figure 2).

AM I AWAKE OR DREAMING?

SECTION FIVE: BORDERLANDS




FIGURE 15.9 • Recording of a correctly signalled lucid dream. Five clear left-right-left-right eye-movement signals are shown in the EOG channel. Typical for REM sleep: EEG channel shows low-voltage mixed frequency and the muscle tone in the EMG channel is very low. The respiration rate and heart rate increase while performing squats in the lucid dreams (after Erlacher and Schredl, 2008, p. 10).

Why does becoming lucid feel like waking up, or becoming more conscious or more 'myself'? An early suggestion was that high levels of cortical activation might be needed to realise it's a dream (LaBerge, 1988). Twenty years later, EEG studies found differences between lucid and non-lucid dreams in the beta frequency band (13–19 Hz, usually associated with waking), with the largest difference in the left parietal lobe, suggesting a link with language and perhaps the ability to understand the words 'I am dreaming' (Holzinger

et al., 2006). As noted earlier, increases in the gamma range have also been found, especially at 40 Hz, as well as increased 'global networking' across the brain. In other words, more long-range connections are made during lucid dreaming, which might mean more links between self-processing, memory, and thinking (Voss et al., 2009).

In REM sleep, the dorsolateral prefrontal cortex (DLPFC), an area whose functions include planning, working memory, and cognitive flexibility, is deactivated compared with waking. If this explains our lack of insight in ordinary dreams, then we might expect DLPFC to be more active in lucid than ordinary dreams, and this was found in the first ever study of lucid dreaming in an fMRI scanner (Dresler et al., 2012). As for the sense of self, the precuneus, on the inner side of the parietal lobes, is also deactivated during REM (Dresler et al., 2012). Since this area relates to self-referential processing, including first-person perspective and the sense of agency, the greater activation of parietal lobes during lucid dreaming might help explain why lucidity brings a sense of being more 'myself'. Although these are only first hints of progress, lucid dreams are no longer considered beyond the pale but are becoming a promising tool for investigating consciousness. We can learn more in this direction by turning our attention to two kinds of experience even stranger than lucid dreams: out-of-body and near-death experiences.

OUT-OF-BODY EXPERIENCES

I was lost in the music, Grateful Dead or Pink Floyd I think it was, and rushing down a dark tunnel of leaves towards the light, when my friend asked 'Where are you?' I struggled to answer, trying to bring myself back to the room where I knew my body was sitting. Suddenly everything became crystal clear. I was looking down on the three of us sitting there. I watched, amazed, as the mouth below said 'I'm on the ceiling'. Later I went travelling, flying above the roofs and out across the sea. Eventually things changed and I became first very small and then very big. I became as big as the whole universe, indeed I *was* the whole universe. There seemed no time, and all space was one. Yet, even then, I was left with the knowledge that 'However far you go, there's always something further'. The whole experience lasted about two hours. It changed my life.

An OBE is an experience in which a person seems to perceive the world from a location outside their physical body. This definition is important because it is neutral as to the explanation required. An OBE is an *experience*, so if you feel as though you have left your body, you have, by definition, had an OBE. During an OBE you feel as though

'you' have left your body and are floating or flying above it, looking down on the world from this new position. You might like to listen to 'And She Was' by Talking Heads, or 'Your Mind Has Left Your Body' by Jefferson Starship for musical evocations. Or you could read Ernest Hemingway's novel *A Farewell to Arms*, which was based on the author's experiences of fighting in Italy in World War I. In Chapter 9 of the novel there is an excellent description of his protagonist having an OBE (as part of a near-death experience) in a dugout during battle: it starts with a flash and a noise that he experiences (synaesthetically) as white and red, and then he can't breathe and feels he is rushing out of himself. He is certain he is dead, realises he'd been wrong to think that death is the end, and then feels himself sliding back into his body and finally alive again on the torn-up ground. We wanted to reproduce the passage here but his publisher forbade us from using it. Like many other descriptions of OBEs, it leaves open for investigation the critical question whether anything leaves the body or not.

OBEs are related to three other types of 'full body illusion', all resulting from displacement of the body schema (Blackmore, 2017). First, 'autoscopy' literally means seeing oneself, but in psychiatry refers to experiences of seeing a double or doppelgänger. The person still seems to be inside their own body but sees an extra self, or a person who looks like them, elsewhere. Second is 'heautoscopy', an even more confusing experience in which people are uncertain whether they identify with their own body or with the double; they may even alternate between one and the other. Finally there is the 'sense of presence' or 'feeling of a presence', a powerful feeling that there is someone else close by even if they cannot be seen. This can happen during sleep paralysis or on the edges of sleep, for example when children are convinced there is a monster under the bed or in the wardrobe. These three, along with OBEs, have in common a doubling of the sense of self (Blanke and Mohr, 2005).

Although an odd experience, OBEs are relatively common, with somewhere between 12% and 20% of people claiming at least one during their lifetime (Blackmore, 2017). More precise estimates are hard to obtain because people often misunderstand survey questions: they may, for example, say 'yes' because they have flying dreams, or 'no' if they think 'real' OBEs require proof of actual travelling. A few people have




FIGURE 15.10 • Self-processing at the temporo-parietal junction (TPJ). The figure summarises data from several neuroimaging studies that revealed an activation of the TPJ during different aspects of self-processing such as visuo-spatial perspective taking, agency, self-other distinction, mental own-body imagery, biological motion perception, and vestibular and multisensory perception. Activations during these paradigms that were found in other areas are not shown. An approximate location with respect to the gyral and sulcal surface is given for each study. Most of the results were found in the right TPJ only or showed a right-hemisphere dominance (Blanke and Arzy, 2005, p. 22).




FIGURE 15.11 • Tunnel of leaves




FIGURE 15.12 • Four types of autoscopic phenomenon. Blue shows the position of the physical body and yellow the phantom or imagined second body. In heautoscopy the experiencer cannot tell which is which.

frequent OBEs, especially during childhood, and even fewer learn to control them.

People have had OBEs when walking down the street, sitting quietly, or even driving a car, and apparently carried on with what they were doing, but the vast majority happen when people are relaxed and lying down. OBEs usually last only a few seconds or minutes, but in rare cases last for many hours. In ‘parasomatic’ OBEs, people seem to inhabit a second duplicate body outside the physical one; in ‘asomatic’ OBEs, they are just a disembodied awareness or a point of consciousness (Green, 1968b; Alvarado and Zingrone, 2015).

OBEs (people who have OBEs) report more psychic experiences and greater belief in the paranormal than others, as well as better dream recall and more frequent lucid dreams (Irwin, 1985; Gackenbach and LaBerge, 1988; Blackmore, 2017). There is no correlation with age, sex, educational level, or religion, nor with standard personality measures, but OBEs do score higher on measures of hypnotisability, capacity for absorption, and positive schizotypy. The concept of schizotypy is based on the idea that schizophrenia lies at one end of a continuum running from normal dissociative and imaginative tendencies to extremely pathological ones. High schizotypes have lots of unusual experiences, disorganised thoughts, flat emotion, and unstable mood and behaviour, but more positively they are also more creative, and there is evidence that OBEs are often ‘healthy schizotypes’ (McCreery and Claridge, 2002), reporting more dissociative experiences and more hallucinations (De Foe, van Doorn, and Symmons, 2012; Parra, 2010).

OBEs have often been dismissed as pathological dissociation, but although in rare cases epilepsy and brain damage can lead to OBEs, the majority are not associated with any pathology. In one study a group of hospitalised schizophrenics reported the same frequency of OBEs as a control group (Blackmore, 1996b), and after studying a very large group of American OBEs, researchers concluded that their ‘psychological health is generally excellent, ranking with the healthiest groups in the population’ (Gabbard and Twemlow, 1984, p. 40).

Precipitating factors include relaxation, reduced sensory input, and vestibular disturbances, as occurs on the verge of sleep. So are OBEs just a special kind of dream? In surveys, OBEs often say that the world looks as real as or even ‘more real’ than usual. Some describe OB vision as brighter and clearer than normal, even claiming a kind of 360 degree vision, but others say it is dim or confusing. In rare cases, time and space seem to disappear as in mystical experiences. OBEs can feel somewhat like lucid dreams in that one feels fully conscious and able to fly around at will, but physiological studies using EEG, heart rate, and other measures show that OBEs induced in the lab occur in a relaxed waking state similar to drowsiness, but not in deep sleep and certainly not in REM sleep (Tart, 1968).

OBEs are not easy to induce, although there are lots of popular books describing how to do it. In the early days of psychical research, hypnosis was used to induce ‘travelling clairvoyance’ or ‘astral projection’, while later experiments tended to use relaxation and imagery exercises. Some drugs can induce OBEs, especially the psychedelics LSD, psilocybin, DMT, and mescaline, but none of these provides

anything like a magic OBE pill. The closest any drug comes to that is probably the dissociative anaesthetic ketamine, which, in sub-anaesthetic doses, paralyses the muscles before inducing unconsciousness. This leads to feelings of body separation and floating but not often to full OBEs. The increased chance of experiencing an OBE after taking drugs of this kind suggests an underlying role for neurotransmitters like dopamine in out-of-body experiences as well as drug-induced 'altered states' and near-death experiences. As well as being a crucial part of the reward system in the brain, dopamine is known to help regulate interpretive tendencies. Dopamine receptors are affected by drugs like LSD (Vollenweider and Kometer, 2010), and dopamine is associated with hallucinatory experiences in diseases like Parkinson's (Fénelon et al., 2000), so there are connections at the level of brain mechanisms between many of these phenomena.

What do OBEs tell us about consciousness? While some people take them as proof that consciousness is independent of the body, there are many other possible explanations.

THEORIES OF THE OBE

OBEs are often so compelling that people become convinced that their consciousness left their body and can survive death, even though neither of these conclusions follows logically from the *experience*. Nineteenth-century psychical researchers thought that the soul or consciousness could be 'exteriorised' during 'travelling clairvoyance', before separating permanently at death. At the same time the new religion of Theosophy, based loosely on a combination of Hindu and Buddhist teachings, taught that we each have multiple bodies: physical, etheric, astral, and several higher bodies. When consciousness leaves the physical as an astral body, sometimes remaining connected by a silver cord, the experience is known as 'astral projection' – a concept that remains popular.

Such theories are forms of dualism and face the same problems (Chapter 1). For example, if the soul or astral body really sees the physical world during projection, then it must be interacting with it and hence it must be a detectable physical entity, yet it is supposed to be non-physical. Many attempts to detect it have been made, including photographing astral bodies, catching them in cloud chambers, or trying to detect them with people, animals, and many types of physical instrument, all to no avail (Morris et al., 1978). On the other hand, if the astral body is non-physical, then it cannot interact with the physical world so as to see it. There are other problems, too. If we can see and hear and remember so clearly with our conscious astral body, why should we need physical eyes, ears, and brain at all?

It is perfectly understandable that OBEs encourage dualist conclusions; they may even explain the origin of the concept of the soul. As Metzinger puts it, 'For anyone who actually had that type of experience it is almost impossible not to become an ontological dualist afterwards' (Metzinger, 2005, p. 78). This is precisely what happened to Sue after an OBE in 1970, and it was only after years of research that she began to change her mind.



FIGURE 15.13 • In the nineteenth century, psychical researchers (almost all male) hypnotised mediums (usually female) to test for 'travelling clairvoyance'. The medium's spirit was supposedly able to travel great distances and report on what it saw there (Carrington, 1919, p. 152).




FIGURE 15.14 • The most obvious theory of OBEs is that a spirit, soul, or astral body leaves the physical and can travel without it. This faces serious problems. What is the phantom made of? How does it communicate with the physical body? Does it travel in the physical world or a replica world of thoughts? How can it gain information from the world without eyes and ears, and without being detected?

'we are talking about the spot where the mind, body, and spirit interact'

(Morse, 1992, p. 211)

'when parts of the TPJ are not working properly, the body schema goes haywire and an OBE results'

(Blackmore, 2017, p. 131)

'The soul is the OBE-PSM'

(Metzinger, 2009, p. 85)

Are these dualist conclusions warranted? Apart from efforts to detect the astral body or soul, some OBEs claim to be able to see events at a distance that they could not possibly have known about unless they were truly 'out-of-body'. Yet despite many popular claims, reliable evidence for this too is lacking. Even the most famous of spontaneous cases tend to crumble on investigation (Blackmore, 2017), and laboratory experiments testing for paranormal perception during the late twentieth century gained only rare hints of success and garnered much controversy (for reviews see Alvarado, 1982; Blackmore, 1982; Irwin, 1985). If such paranormal claims could be verified, they would dramatically change our understanding of OBEs and potentially of consciousness too, but the evidence so far is weak and there have been no more recent experiments of this kind.

OBEs, CONSCIOUSNESS, AND NEUROSCIENCE

The alternative to astral projection or any other dualist theory is to say that, despite how it feels, nothing actually leaves the body. Among theories of this kind, early psychoanalytic theories described the OBE as a dramatisation of the fear of death, regression of the ego, or reliving the trauma of birth, and Jung saw it as part of the process of individuation. But such theories are largely untestable and have led to no advances in our understanding.

Early psychological theories generally started from the finding that OBEs occur when sensory input is reduced or disrupted, proposing different responses to this disruption (Palmer, 1978; Irwin, 1985; Blackmore, 2009, 2017). For example, the cognitive system might try to construct a new (if inaccurate) body image and a new 'model of reality' derived from memory and imagination, using the kinds of bird's-eye views that are common in memory and dreams. This was supported by evidence that OBEs are better at spatial imagery and at switching viewpoints in imagery, and more often dream in bird's-eye view (Blackmore, 1996b). Once the necessary neuroscience was available, the peculiarities of the OBE began to fall into place.

The temporal lobe has long been implicated in OBEs because temporal-lobe epileptics report more OBEs as well as psychic and mystical experiences. The Canadian neuroscientist Michael Persinger (1983, 1999) proposed that all religious and mystical experiences are artefacts of temporal lobe function. He has succeeded in inducing OBEs, body distortions, the sense of presence, and many other experiences using his own version of TMS, with stimulation on the left side producing a sense of presence and on the right side OBEs.

An early hint of a more precise connection was found accidentally in the 1930s when American neurosurgeon Wilder Penfield electrically stimulated the brain of an epileptic woman, trying to find the seizure focus. On one occasion, when stimulating her right temporal lobe, she cried out 'Oh God! I am leaving my body' (Penfield, 1955, p. 458).

Over half a century later, with much finer electrodes and greater precision, a team of neurosurgeons in Geneva achieved the same result with another epileptic patient.

When a weak current was passed through a subdural electrode on the right angular gyrus, she reported sinking into the bed or falling from a height. With increased current she said, 'I see myself lying in bed, from above, but I only see my legs and lower trunk'. This experience was induced twice more, as were various body image distortions. The researchers attributed her OBE to a failure to integrate somatosensory and vestibular information caused by the stimulation (Blanke et al., 2002).

The specific area involved was the right temporoparietal junction (TPJ). In this area visual, tactile, proprioceptive, and vestibular information all come together to construct a body schema. This is the bodily representation that is needed by all animals and is constantly updated as we move about. It underlies our physical or bodily sense of self and is integrated with emotions and memories, and with ideas about how we appear to others, which all in turn contribute to generating our body image and self image.

Several lines of research have converged to show how the OBE relates to self-processing at the TPJ. Not only does direct stimulation of this spot induce OBEs, but PET scanning has shown brain activation at the TPJ during OBEs induced by stimulating the right temporal gyrus. The researchers conclude that 'activation of these regions is the neural correlate of the disembodiment that is part of the out-of-body experience' (de Ridder et al., 2007, p. 1829). Other evidence comes from several patients who experience OBEs or autoscopy and have been found to have damage to the TPJ (Blanke et al., 2004; Blanke and Arzy, 2005).

The right TPJ is also involved in perspective-taking, the ability to see things from another's point of view. A visual test is the Own Body Transformation Task, which entails looking at rotated human figures and deciding which is their right hand. Evoked potential mapping shows selective activation of the TPJ during this test, and interfering with the TPJ using TMS makes this mental transformation more difficult (Blanke and Arzy, 2005). Yet there is much doubt here.

If OBEs depend on 'disturbed processing at the TPJ' or 'disruption of vestibular-motor integration' (Wilkins, Girard, and Cheyne, 2011), then OBEs might be expected to do worse at perspective-taking. But we noted earlier some evidence of the opposite, and British psychologist Jason Braithwaite has argued that rather than implying failure, 'genuine OBEs should not be regarded as a flaw in the system of certain individuals but as "the other side of the coin" of full-blown perspective taking' (Kessler and Braithwaite, 2016, p. 423). In this case, OBEs might do better at tasks involving perspective-taking, as they did in early experiments on changing viewpoints (Blackmore, 1996b). Jason Braithwaite and colleagues (2013) devised an improved version of the OBT and found that OBEs did perform better. Whether OBEs reveal a skill or a flaw is very much an open question.



FIGURE 15.15 • Cloud chambers are normally used to detect subatomic particles, which leave a trail of water droplets as they pass through super-cooled water vapour. In the 1920s physical researchers adapted the technique to detect the astral double. A frog, mouse, or grasshopper was placed in the inner chamber and poison introduced to kill the animal whose soul would then pass out through the cloud chamber and so be made visible.



FIGURE 15.16 • This picture of a 'phantom frog', kept in the SPR archives, was taken in the 1930s by R.A. Watters, who claimed to have photographed the 'intra-atomic quantity' departing from the physical body at death.

OUT-OF-THE-BODY IN VIRTUAL REALITY



FIGURE 15.17 • An alternative to laboratory tests. For several years, during the 1980s, Sue displayed targets in the kitchen, out of view of the window, so that anyone who claimed to have OBEs could try to see them. These were a five-digit number, one of twenty common words, and one of twenty small objects. They were selected using random number tables, and changed regularly. OBEs could try to visit from their own home, or anywhere else, during spontaneous OBEs, but none successfully reported the targets.

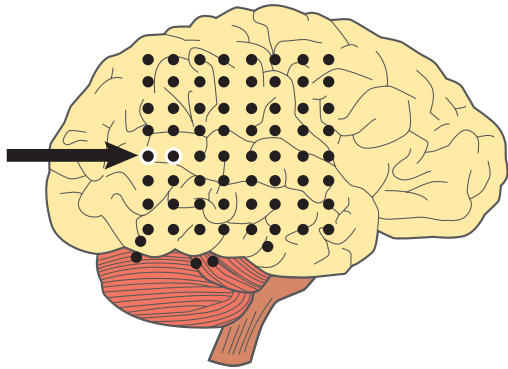


FIGURE 15.18 • 3D surface reconstruction of the right hemisphere of the brain from magnetic resonance imaging. Subdural electrodes were implanted in the brain of an epileptic patient undergoing presurgical evaluation; the locations at which focal electrical stimulation (ES) evoked behavioural responses are shown. Out-of-body experiences (OBEs), body-part illusions, and vestibular responses were induced at the site marked with the arrow (Blanke et al., 2002, p. 269).

body in front of him and wanting to 'slip into' it. In further experiments volunteers were shown either their own back, the back of a fake mannequin, or a large wooden slab. With synchronous stroking of either their own back or the mannequin many felt as though the virtual body was their own, and some felt they could 'jump into' it.

In a different method, in Sweden Henrik Ehrsson (2007) also used head-mounted displays showing the volunteer's own back, but in this case they stroked the person's chest while moving a rod up and down in front of the cameras to make it look as it would to someone having their chest stroked. With this set-up the volunteers reported that they seemed to move backwards towards the position of the camera, not forwards.

Both methods have since been combined (Lenggenhager, Mouthon, and Blanke, 2009) and even tested with volunteers lying in an fMRI scanner (Ionta et al., 2011), revealing, once again, a central role for the TPJ. Among other findings are that when these illusions are successfully induced, body temperature drops, pain is felt less strongly, and when threatened with a knife people react less when the illusion is strongest (Guterstam and Ehrsson, 2012). This research has moved extraordinarily fast and taken the OBE from a topic shunned by mainstream psychology to one that is actively being investigated for what it can tell us about just how and why we normally build a self-model that coincides with our bodily position, and occasionally build one that seems to fly.



PRACTICE 15.3 WHAT SURVIVES?

As many times as you can, every day, take a good look at your own body and ask yourself '**When this body is gone, what will remain?**' Try to strip away everything that you know will turn to ashes or dust and then imagine or think or feel what might be left.




FIGURE 15.19 • Two methods for inducing out-of-body illusions using virtual reality. In both the participants wear a head-mounted display showing images from cameras two metres behind them. In Lenggenhager's (2007) method (a), the participants can see their own back being stroked, and get the sense of moving forwards. In Ehrsson's (2007) method (b), they feel their chest being stroked while seeing a stick appearing and disappearing in front of the camera, and get the sense of moving backwards.

Overall, OBEs give us insights into how our sense of self is normally constructed and what happens when the normal mechanisms for anchoring our sense of self in our bodies temporarily break down. They remind us from another angle that the sense I have of being in my head looking out through my eyes is not a truth, but just the result of all the neural and other processes that generally make 'in my head' feel like where 'I' am. The switch of that location in an OBE makes us realise it is always just a construct – but also shows just how profoundly consciousness changes when it no longer seems to be grounded in the body.

NEAR-DEATH EXPERIENCES

Across many ages and cultures, people coming close to death report a consistent set of experiences, including the 'returned from the dead' writings in Tibetan Buddhism, a description in Plato's *Republic*, and myths from ancient Greece, Native Americans, and contemporary European folklore. Nineteenth-century psychical researchers collected accounts of 'death-bed visions' reported by people just

before they died, but as medical expertise grew in the twentieth century, 'near-death experiences' (NDEs) were reported by survivors of cardiac arrest.

The term 'NDE' was coined in 1975 by American physician Raymond Moody, who interviewed about fifty survivors of close brushes with death and produced a composite account (see the website for more detail). Subsequent studies broadly confirmed the main components: a tunnel, an OBE, a brilliant white or golden light, positive and loving emotions, visions of another world, meetings with other beings, a life review, and the decision to return (Ring, 1980; Murray, 2009). NDEs cannot be attributed to medication given near death because they tend to be less complex, not more so, with medication. Nor can they be explained solely by lack of oxygen. In some cases, blood oxygen levels have even been found to be higher in NDErs than in other patients (Parnia et al., 2001). Most important is that NDEs also occur in people who are far from death, such as those in dangerous experiences who think they are going to die, like Hemingway's soldier surviving shell fire, or a mountaineer escaping unhurt from a terrifying fall.

Most NDEs are pleasant and even blissful, but rarer hellish experiences include black voids and nothingness, chattering demons, black pits, naked zombie-like creatures, and other symbols of traditional hell. On some estimates up to 15% of NDEs are hellish, but it is hard to be accurate because people may be keen to forget them and unwilling to talk about them. Interestingly, suicide attempters generally report positive NDEs and are less likely to try to kill themselves again. Highly positive after-effects are common, including greater interest in spirituality and in caring for others, and reduced interest in material belongings or success. These effects can be long-lasting, with NDErs in one study still reporting continued positive changes eight years after their brush with death (van Lommel et al., 2001). Less often NDErs are left depressed and a few find themselves estranged from family and friends by the changes that take place.

The early studies collected accounts retrospectively, making it impossible to know how common NDEs are, but later prospective studies found out. In Britain, medical researcher Sam Parnia and his colleagues (Parnia et al., 2001) interviewed all survivors of cardiac arrest in a Southampton hospital during one year. Seven out of sixty-three (11%) reported memories, of which four counted as NDEs on the Greyson NDE scale. None had an OBE.

In the USA, a thirty-month study of 1595 consecutive patients admitted to a cardiac care unit found that among those who suffered a cardiac arrest, 10% reported NDEs compared with 1% of other patients (Greyson, 2003). Further prospective studies have found the incidence of NDEs in survivors of cardiac arrest to be between 9% and 23% (for a review see Blackmore, 2017, pp. 241–242).

The most-cited study of this kind was by cardiologist Pim van Lommel and his colleagues in the Netherlands. They studied 344 consecutive patients resuscitated after cardiac arrest. Sixty-two (18%) reported some memories and forty-one (12%) described a core experience (including out-of-body, tunnel, and light experiences), but NDEs did not depend on the duration of cardiac arrest or medication received. Thirty-seven of the NDErs were interviewed two years later and nearly all retold their experiences almost exactly. When compared with those who had no NDE, they had increased belief in an afterlife, less fear of death, a greater interest

'at the time of physical death consciousness will continue to be experienced in another dimension, in an invisible and immaterial world'

(van Lommel, 2006, p. 148)

in spirituality, and increased love and acceptance for others. Eight years after the events, all the patients claimed positive changes (van Lommel et al., 2001).

INTERPRETING NDES

Dismissing NDEs as fabrications or wish fulfilment is unreasonable. The similarities across ages and cultures, and the reliability of the findings, suggest that NDEs have something interesting to teach us about death and consciousness. The question is, what?

A common reaction, as to OBEs, is that NDEs are proof of dualism – of the existence of a soul or consciousness that operates independently of the brain and can survive death. For Kenneth Ring (1980), the experiences ‘point to a higher spiritual world’ and access to a ‘holographic reality’; for Parnia and Fenwick (2002), understanding NDEs will require ‘a new science of consciousness’; for van Lommel (2009), they are evidence for non-local consciousness or ‘endless consciousness’.

Two types of evidence are commonly given in support. First, NDErs describe ‘clear’ states of consciousness with lucid reasoning and memory when their brain is severely impaired. ‘How could a clear consciousness outside one’s body be experienced at the moment that the brain no longer functions during a period of clinical death with flat EEG?’, ask van Lommel and colleagues (2001, p. 2044). Indeed, how could it? If ‘clear consciousness’ were really possible with no heart-beat and a completely flat EEG, this would indeed change our view of the mind–brain relationship, but this has not been demonstrated. The problem concerns timing. There is not one case in which we know that the experiences occurred when the person’s brain was not functioning; the NDEs could just as well have occurred just before, during, or just after the medical crisis. A number of physiological differences have been found between NDErs and control groups, but arguments continue about what this means for our understanding of mind, brain, and consciousness (Trent-Von Haesler and Beauregard, 2013), and in particular it is unclear whether the changes are a cause or a consequence of the NDE.

Second, there are many claims of the paranormal, including compelling accounts of people seeing things at a distance which they could not possibly have known about. Yet these cases have not stood up well to investigation (for a review see Blackmore, 2017). For example, van Lommel supports his claims of ‘endless consciousness’ and ‘memory outside the brain’ (2013) with a decades-old anecdote reported to him second-hand about someone commonly known as ‘dentures man’ and which even believers in life-after-death have concluded is unconvincing (Smit, 2008).

One way to find out whether consciousness persists beyond physical death would be to provide randomly selected, concealed targets that NDErs could see during their experience. The best study of this kind is AWARE (AWAREness during REsuscitation), a multi-hospital project launched in 2008 to measure brain function at the same time as providing hidden images that NDErs might be able to see. Sadly, none of the patients who had NDEs looked at the hidden targets (Parnia et al., 2014). One man did have an OBE and this occurred around 20 or 30 seconds into his three-minute cardiac arrest. Interestingly, odd bursts of activity have

‘An individual who should survive his physical death is also beyond my comprehension [. . .]; such notions are for the fears or absurd egoism of feeble souls’

(Einstein, 1949/2006, p. 7)



FIGURE 15.20 • In Victorian times most people died at home, surrounded by their families. Reports of death-bed experiences were common, including other worlds, beautiful music, and visions of those who had already ‘passed over’ coming to greet the newcomer. Occasionally observers said they saw the dying person’s spirit leaving the body and going up into the light (Muldoon and Carrington, 1929, p. 186; according to clairvoyant vision of Andrew Jackson Davis).

WHEN THIS BODY IS GONE, WHAT WILL BE LEFT?

‘NDEs tell us nothing about life after death’

(Blackmore, 1993, p. 4)

previously been recorded at about this time in dying patients, and studies with rats show a similar burst of activity, probably due to cortical disinhibition, 20 to 30 seconds after their hearts stop (Chawla et al. 2009). This underlines again the importance of finding out just when NDEs are occurring before jumping to conclusions about consciousness beyond death.

Van Lommel’s research itself was impressive, but his conclusions do not follow from the findings. Braithwaite concludes that ‘Despite its impact in NDE circles, the van Lommel et al. study provides no evidence that human consciousness survives bodily death’, and ‘poses no serious challenge at all to current neuroscientific accounts of the NDE’ (Braithwaite, 2008, p. 15).

An alternative, naturalistic approach to understanding NDEs depends on the findings that all the components of the classical NDE can be caused by cortical disinhibition and excessive uncontrolled brain activity. This can occur in conditions of severe stress, extreme fear, and cerebral anoxia, as well as with certain drugs, and we already have most of the ideas needed to understand why this should cause NDEs. Tunnels and lights are caused by disinhibition in visual cortex, and strange noises by disinhibition in auditory cortex. OBEs and life reviews can be induced by heightened temporal lobe activity, and the positive emotions and lack of pain have been attributed to the action of endorphins and enkephalins, endogenous opiates which are widely distributed in the limbic system and released under

stress. The visions of other worlds and spiritual beings might be real glimpses into another world, but against that hypothesis is that the worlds described tend to fit people’s cultural upbringing and religious beliefs. In the popular genre of ‘Heaven tourism’, Christians report seeing Jesus, angels, and a door or gate into heaven. Yet Hindus are more likely to meet the king of the dead and his messengers, the Yamdoots.

All these apparently strange experiences – sleep paralysis, lucid dreams, OBEs, and NDEs – once seemed inexplicable. But now that we are beginning to understand them, they seem not to provide evidence for other worlds or consciousness beyond the brain but, like dreaming or psychedelic states, to offer important test cases for intuitions about the relations between conscious and unconscious, between real and unreal, between self and body, and between retrospective verbal report and ‘experience itself’. Not least, these times when sensory input and bodily interaction with the world are reduced and experience seems more self-generated encourage us to reflect on who or what is doing the generating – that is, on the nature of our selves. This is one more version of the question that has been with us throughout the book, in our explorations of everything from Cartesian audiences and ghosts in machines to free will and social robots, and which we tackle head-on in the next chapter: who is the one who has conscious experiences? Who or what are you?



Blackmore, S. (2017). Incredible! A chapter in S. Blackmore, *Seeing myself: The new science of out-of-body experiences* (pp. 276–292). London: Robinson.

Assesses the evidence for paranormal events during NDEs, critiquing van Lommel’s claims of ‘Endless consciousness’. Compare this with van Lommel et al. (2001), and see also the previous two chapters (pp. 235–275) for an overview of NDEs.

Hobson, J. A., Pace-Schott, E. F., and Stickgold, R. (2000). Dreaming and the brain: Toward a cognitive neuroscience of conscious states. *Behavioral and Brain Sciences*, 23(6), 793–1035.

Reviews evidence for correlations between the phenomenology and the physiology of dreams, and introduces the AIM model. Commentaries (pp. 843–1018) include Stephen LaBerge’s on lucid dreaming (pp. 962–963, with figures at www.lucidity.com/slbbs/index.html).

Metzinger, T. (Ed.) (2009). *The ego tunnel*. New York: Basic Books, pp. 82–101.


Uses his own and others’ experiences of OBEs to ask what is needed for the fundamental level of self: minimal phenomenal selfhood.

Nir, Y., and Tononi, G. (2010). Dreaming and the brain: From phenomenology to neurophysiology. *Trends in Cognitive Science*, 14(2), 88.

The differences between waking and dreaming consciousness, and how multidisciplinary study helps us relate dreams to visual imagery.

van Lommel, P., van Wees, R., Meyers, V., and Elferich, I. (2001). Near-death experience in survivors of cardiac arrest: A prospective study in the Netherlands. *The Lancet*, 358(9298), 2039–2045.

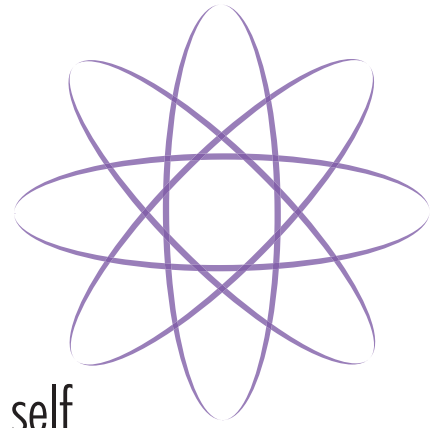
Includes follow-up interviews and the claim that NDEs provide evidence for consciousness beyond the brain.



Self and other

SECTION

Six



Egos, bundles, and theories of self

CHAPTER

SIXTEEN

Who is reading this book? Who is conscious of the writing on the page, the attempt to understand and answer the question, or the sounds of revelry in the next room?

Questions about the nature of consciousness are intimately bound up with those about the nature of self because it *seems* as though there must be someone having the experience: that there cannot be experiences without an experiencer. Our experiencing self seems to be at the centre of everything we are aware of at a given time, and to be continuous from one moment to the next. In other words, it seems to have both unity and continuity. In Chapter 6, we explored some reasons for questioning the idea that conscious experience is unified, but we might still be tempted to attribute unity and continuity to the self who has the experiences. More problems arise, however, when you ask what kind of thing that experiencer might be.

In everyday language, we talk unproblematically about our 'self'. 'I' got up this morning, 'I' like muesli for breakfast, 'I' can hear the robin singing, 'I' am an easy-going sort of person, 'I' remember meeting you last week, 'I' want to be an engine driver when I grow up. 'I' distinguish 'myself' from 'you' and 'yourself'. It seems that we not only think of this self as a single thing, but accord it all sorts of attributes and capabilities. In ordinary usage, the self is the subject of our experiences, an agent who carries out actions and makes decisions, a unique personality, and the source of desires, opinions, hopes, and fears. This self is 'me'; it is the reason why anything matters in 'my' life. But where or what is this 'me'?

'An experience is impossible without an experiencer'

(Frege, 1918/1967, p. 27; in Strawson, 2006, pp. 189–190)

'I am conscious that I exist, and I who know that I exist enquire into what I am'

(Descartes, 1641/2008, p. 81)

Every life is in many days, day after day. We walk through ourselves, meeting robbers, ghosts, giants, old men, young men, wives, widows, brothers-in-love, but always meeting ourselves.

(James Joyce, *Ulysses*, 1922)

One way of escaping the problem might be to declare that I am my whole body, and there is no need for a self as well. This would be fine, except that most people don't feel that way. The 'whole body' idea of self works for some purposes – 'I' went shopping, 'I' tripped on the carpet, 'I' am an expert skier; 'she' went on holiday and 'he' popped in for a drink. But it works less well for others. Can you really say that your whole body believes in eliminative materialism, is worried about your parent's health, or hopes it won't rain tomorrow?

Hold up your hand in front of you. Whose hand is this? Look down at your feet. Whose feet are these? Perhaps you feel as though you and the hands and feet are one, and there is no gap between a perceived and a perceiving self. Or perhaps you feel as though the hand is over there and 'you' are in here, somewhere behind 'your' eyes and looking out at it, this thing that belongs to you. In this

case, who are 'you'? Who is calling this 'my hand', 'my body', and even 'my brain'? In this and many other ways, we come to feel as though we are not the same as our body but are, to use an old traditional metaphor, something like the driver of a carriage or the pilot of a ship. I talk about the body as something that 'I' possess. And so I separate 'myself' from it.

That this apparently natural way of thinking about ourselves is problematic has been recognised for millennia. In the sixth century BC, the Buddha challenged contemporary thinking with his doctrine of *anatta*. This is often, perhaps inaccurately, translated as 'no-self', when really he was rejecting the common idea that we consist of a separate and continuous entity. Instead, he claimed that the self is just a name or label given to a collection of parts, in the way that we give the name 'carriage' to a set of parts – a suggestion that seems as hard to understand and accept today as it was then. The Greek philosophers struggled with similar issues, including Plato, who wanted to know whether the psyche (the soul or true essence of a person) is immortal. In his famous dialogues, he argued both that the psyche is immortal and that it has parts – appetitive, emotional, and rational parts. This created a serious problem, since he also believed that only a unitary and indivisible thing could be immortal. His ideas

helped set the stage for millennia of mind–body dualism in which the supposedly rational immaterial self has been consistently valued over the supposedly base and animal-like bodily self.



FIGURE 16.1 • The Buddha taught the doctrine of *anatta* or 'no-self'. Parfit (1987, p. 21) calls him the first bundle theorist.

Similar problems have plagued many thinkers since. In philosophy, there are numerous theories of the nature of self (or what persons are), of personal identity (or what makes someone the same person over time), and of moral responsibility. In psychology, researchers have studied the development of the sense of self in children, the construction of social selves, self-attribution, the factors affecting personal identity, dissociative states, and various pathologies of selfhood. We cannot consider all of these here, so in this chapter we will concentrate on a few of those theories that are most relevant to consciousness.

The central question is why it seems as though I am a single, continuous self who has conscious experiences. Possible answers can be divided into two major types. The first claims that it is true: there really is some kind of continuous self that is the subject of my experiences, that makes my decisions, and so on. The second accepts that it seems this way but claims that really there is no underlying continuous and unitary self. The illusion that there is has to be explained some other way. Oxford philosopher Derek Parfit (1984, 1987) described these two types as ‘ego theories’ and ‘bundle theories’, citing the Buddha as the first bundle theorist.

Ego theories are undoubtedly the more popular. Many religions entail notions of spirits or souls, including both Christianity and Islam, which teach that the soul is a continuing entity that is central to a person’s life, underlies moral responsibility, and can survive the death of the physical body. Among the major religions, Buddhism alone denies the existence of such entities.

Perhaps this is not surprising when we learn that young children seem to be natural dualists. According to psychologist Paul Bloom, children as young as three see the world as containing two distinct domains, bodies and souls. By five or six they may have learned that the brain does lots of useful things, like thinking and solving problems, but they still imagine it as ‘a tool we use for certain mental operations [. . .] a cognitive prosthesis, added to the soul’ (Bloom, 2004, p. 201). Other research found that when 5–6-year-olds were given a hamster that had apparently been duplicated by a very special machine, they thought that fewer of the original’s episodic memories than its physical properties would be transferred to the duplicate hamster (Hood et al., 2012). Interviewed about the functioning of dead agents, children as young as four know that biological functions (such as needing the toilet or food) stop at death, but they make no distinction between psychobiological or perceptual states on the one hand and epistemic or emotional states on the other, often believing that both continue after death. As they get older, they are more likely to separate these out, with older children and adults attributing beliefs, emotions, and desires to the dead, but not perceptions: ‘default “afterlife” beliefs are pruned in a systematic fashion during development’ (Bering and Bjorklund, 2004, p. 229). One interpretation is that we attribute to the dead those mental states that we cannot imagine being without (Bering, 2002). This may be one reason why ego theories are so prevalent and hard to shift: because we cannot imagine being without what feels like a conscious self, we may be tempted to grant it continuity beyond death as well as in life.

Most forms of substance dualism are ego theories because they equate the separate mind, or non-physical substance, with the experiencing self. One example is Popper and Eccles’s dualist interactionism (Chapter 6), in which the self-conscious mind controls its brain and scans the brain’s activity. But the distinction between

‘a Bundle Theorist denies the existence of persons . . . Bundle Theory is hard to believe’

(Parfit, 1987, pp. 20, 23)



PROFILE 16.1

David Hume (1711–1776)



David Hume was born in Edinburgh and studied law at Edinburgh University, although he never graduated. He tried his hand at commerce in Bristol, but nearly had a nervous breakdown. In 1734 he moved to France and there wrote his masterpiece, *A Treatise of Human Nature*, in his mid-twenties. This long book was not a great success, but the shortened version, *An Enquiry Concerning Human Understanding*, became a classic. He built on the empiricism founded by Locke and Berkeley and wrote on causation, morals, and the existence of God. Hume distinguished between ‘ideas’ and ‘impressions’ according to the force and liveliness with which they make their way into consciousness. He reported that he could never catch himself without a perception, and never found anything but the perceptions, which is why he concluded that the self is not an entity but a ‘bundle of sensations’.

‘I can never catch myself’

(Hume, 1739, Section VI)

‘I am not thought, I am not action, I am not feeling: I am something that thinks, and acts, and suffers’

(Reid, 1785, p. 318)

ego and bundle theories should not be confused with the distinction between dualism and monism or materialism. As we shall see, many materialist scientists, while denying dualism, do believe in a persisting self.

Bundle theories take their name from the philosophy of David Hume, who argued in *A Treatise of Human Nature* (1739) that we are nothing but a bundle or collection of different perceptions, which succeed each other with an inconceivable rapidity, and are in a perpetual flux and movement. All our sensations, impressions, and ideas seem to be tied together because memory gives them apparent continuity, and as such is the source of personal identity. There is no additional unified entity that experiences things or holds the experiences together. He wrote:

For my part, when I enter most intimately into what I call *myself*, I always stumble on some particular perception or other, of heat or cold, light or shade, love or hatred, pain or pleasure. I never can catch *myself* at any time without a perception, and never can observe any thing but the perception.

(1739, Section VI)

By staring deep into his own experience, Hume, like the Buddha, seems to have discovered that there is no experiencer. Not surprisingly, Hume’s ideas were unpopular, and his denial of self was countered by the commonsense approach of his fellow Scottish philosopher Thomas Reid, who protested: ‘I am not thought, I am not action, I am not feeling: I am something



PRACTICE 16.1

WHO IS CONSCIOUS NOW?

As many times as you can, every day, ask yourself ‘Am I conscious now?’ You will probably be sure that you are, whether you are aware of walking along the road, the room around you, or the music you are listening to. Now turn your attention to whoever or whatever is having this experience. This is presumably what Hume was doing when he made his famous realisation about self. Can you see or feel or hear the *experiencer*, as opposed to the experienced world? At first, you will probably be sure that there is an experiencer, but it may be difficult to see any further. Keep looking. Keep asking ‘**Who is conscious now?**’

This is not an easy exercise, but it will repay practising over many weeks or months. Try to see whether there is a separation between the experienced and the experiencer, and if so, what the experiencer is like. This practice forms the basis of the next exercise as well.

that thinks, and acts, and suffers' (1785, p. 318). The thoughts and actions and feelings may come and go, but the *self* or *I* to which they belong is permanent. In other words, Reid appealed to ego theory.

These two views capture a fundamental split in the way people think about the nature of self. On the one hand, ego theorists believe in continuously existing selves who are subjects of experience and who think, act, and feel. On the other hand, bundle theorists deny there is any such thing.

As Hume knew all too well, bundle theory is counter-intuitive, for the non-existence of my self is difficult even to contemplate. But there are many good reasons at least to try. We will begin with some extraordinary case histories challenging the natural assumption that each human being has one conscious self.

MULTIPLE PERSONALITY

On 17 January 1887, an itinerant preacher called Ansel Bourne walked into a bank in Providence, Rhode Island, withdrew \$551, paid some bills, and got into a horse-car bound for Pawtucket. Nothing more was heard of him for two months. The local papers advertised him as missing and the police hunted in vain.

Two weeks later, a Mr A. J. Brown rented a small shop in Norristown, Pennsylvania, stocked it with stationery, confectionery, and fruit, and set up a quiet trade. He went to Philadelphia to replenish his stock, slept and cooked in the back room, regularly attended church, and, according to neighbours, was quiet, orderly, and 'in no way queer'. Then, at 5 a.m. on 14 March he was woken by an explosion to find himself feeling weak and afraid and in an unfamiliar bed. Calling for help, he said his name was Ansel Bourne, he knew nothing of Norristown or shopkeeping, and the last thing he remembered was taking money out of a bank in Providence. His neighbours thought him insane and so, at first, did the doctor. But, happily, they did as he asked and telegraphed his nephew in Providence. A reply came swiftly back and soon the Rev. Ansel Bourne was taken home.

Early in 1890 William James and Richard Hodgson conceived the idea of hypnotising Bourne to see



CONCEPT 16.1 EGO AND BUNDLE THEORIES OF SELF

Ego theory

The reason each of us feels like a continuous, unified self is because we are. Underlying the ever-changing experiences of our lives, there is a self who experiences all these different things. This self may (indeed must) change gradually as life goes on, but it is still essentially the same 'me'. In other words, according to any kind of ego theory, the self is a continuous entity that is the subject of a person's experiences and the author of their actions and decisions.

Ego theories include:

- Cartesian dualism
- Immortal souls
- Reincarnating spirits
- Gazzaniga's interpreter
- MacKay's self-supervisory system
- Add your own examples . . .

Bundle theory

The feeling that each of us is a continuous, unified self is an illusion. There is no such self, but only a series of experiences linked loosely together in various ways. Bundle theory does not deny that each of us *seems* to be a continuous conscious being. It denies that there is any continuously existing entity which explains that appearance. There are experiences, but there is no one who has them. Actions and decisions happen but not because there is someone who acts and decides.

Bundle theories include:

- The Buddhist notion of *anatta*, or no-self
- Hume's bundle of sensations
- Self as a product of discourse
- Dennett's no audience in the Cartesian theatre
- Add your own examples . . .



FIGURE 16.2 • Dr Jekyll and Mr Hyde: good doctor and evil murderer sharing the same body, from Robert Louis Stevenson's classic 1886 novel, here in the film adaptation directed by Rouben Mamoulian and starring Fredric March (1931).

if they could contact the dissociated personality. When James put Bourne into a hypnotic trance, Mr Brown reappeared, describing the places he had stayed and seeming unaware of any connection with Bourne's life. James and Hodgson tried in vain to reunite the two personalities and Hodgson concluded that 'Mr. Bourne's skull to-day still covers two distinct personal selves' (James, 1890, i, p. 392).

What does this extraordinary case of 'fugue' tell us? At the time, doctors, psychologists, and psychological researchers argued over whether it could be explained by epilepsy, fraud, split personality, psychic phenomena, or even spirit possession (James, 1890; Hodgson, 1891; Myers, 1903). Bourne had blackouts and seizures that might indicate epilepsy, but they could not, on their own, explain the extraordinary phenomena. Perhaps the most obvious thing to note is the connection between memory and selfhood. When the character of Brown reappeared, the memories of that missing time came back and the rest of life seemed vague or non-existent. When Bourne reappeared, the memories of Mr Brown and the whole of his short and simple life were gone. As far as we know Mr Brown never returned, and by late 1887 this personality was gradually disintegrating.

At about that time, Robert Louis Stevenson's fantastic tale of *The Strange Case of Dr Jekyll and Mr Hyde* (1886) was published. By then many real-life cases of what became known as multiple personality had appeared. Hypnosis, or mesmerism, was popular for treating such conditions as hysteria, and occasionally doctors or psychiatrists found that hypnotised patients manifested a completely different personality. These patients, almost always women, did not just reveal different personality traits (the way we use the term 'personality' today), but appeared to be two or more distinct people inhabiting a single body (what we might call persons or selves).

Early in 1898, the Boston neurologist Dr Morton Prince was consulted by Miss Christine Beauchamp (Prince, 1906). She had endured a miserable and abusive childhood and was suffering from pain, fatigue, nervousness, and other symptoms which he treated with both conventional methods and hypnosis. Under hypnosis a second, rather passive, personality appeared (labelled BII), but one day Miss Beauchamp began speaking about herself as 'she' and a third personality called Sally had appeared (BIII). Sally was childish, selfish, playful, and naughty, while Miss Beauchamp was religious, upright, reserved, and self-controlled; Sally was fit and strong while Miss Beauchamp was weak and nervous. During many years of treatment, several more personalities appeared with different tastes, preferences, skills, and even states of health.

Sally used to delight in tricking Miss Beauchamp by taking a long walk in the dark and then 'folding herself up' to leave poor Miss Beauchamp to walk home, terrified and ill. Even worse, Sally tore up Miss Beauchamp's letters, shocked

her friends, and made her smoke cigarettes, which she hated. As Prince put it, she ‘indulges tastes which a moment before would have been abhorrent to her ideals, and undoes or destroys what she had just laboriously planned and arranged’ (1906, p. 2). Some nights Sally threw off all the bedclothes and piled the furniture on the bed before folding up again. Imagine waking up in such a situation, with no recall of the past few hours, and knowing that no one else could have entered your room.

Two of the personalities had no knowledge of each other, or of the third, and each had blanks in memory corresponding to times when the others were active. Oddly enough, though, Sally knew about the others and said she could recall times when they were in control. She even claimed that, though ‘squeezed’ when Miss Beauchamp was ‘out’, she was still conscious and had her own thoughts, perceptions, and will. She claimed to be aware of Miss Beauchamp’s dreams, although she herself neither slept nor dreamt. In other words, this was not alternating consciousness (as we might interpret Ansel Bourne’s case) but simultaneous consciousness, or ‘co-consciousness’, with what Prince calls a ‘subconscious self’ or a ‘subconsciousness’, having its own stream of conscious experiences while another controls the body.

Seeking the ‘real Miss Beauchamp’, Prince concluded that the sub-personalities, including Sally, were just ‘a dissociated group of conscious states’ (1906, p. 234) deserving ‘psychical murder’ (p. 248). This extraordinary story had a happy ending, for he eventually brought them all together into what he called (though others might disagree) ‘the real, original or normal self, the self that was born and which she was intended by nature to be’ (p. 1).

This was among the last of the classic cases of multiple personality – most being reported between 1840 and 1910 before a wave of reaction against the increasingly bizarre phenomena reported. Critics pointed out that most cases involved older men hypnotising young women who were eager to please and might easily be duped. Others argued that multiple personality, or ‘dissociation’, was iatrogenic (i.e. created by the treatment or the therapist), and in 1994, the American Psychiatric Association’s Diagnostic and Statistical Manual of Mental Disorders-IV changed Multiple Personality Disorder (MPD) to Dissociative Identity Disorder (DID). The fifth edition of the DSM (American Psychiatric Association, 2013) defines DID as involving ‘Two or more distinct identities or personality states [. . .], each with its own relatively enduring pattern of perceiving, relating to and thinking about the environment and self’. DID research suggests that cultural context is related both to the overall prevalence and to the likelihood of different triggers, such as experience of abuse (Slogar, 2011).

Should we conclude that multiple personality tells us more about the interactions between patients, hypnotists, and therapists than about self and consciousness?

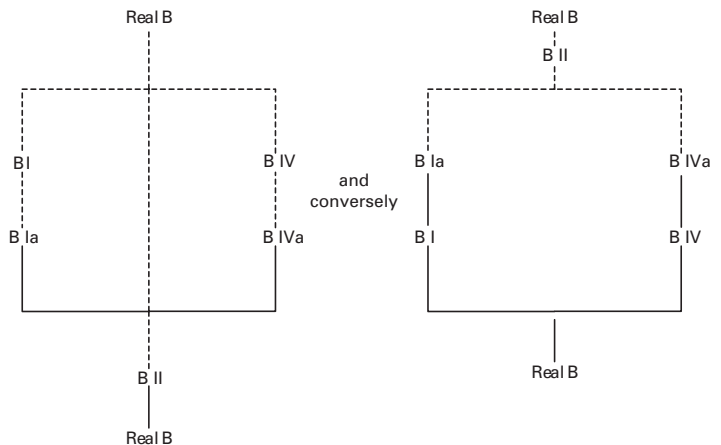


FIGURE 16.3 • According to Morton Prince, the Real Miss Beauchamp disintegrated into BI and BIV. (a) These two personalities could be hypnotised to give BIa and BIVa, who could be synthesised into BII. (b) Alternatively Real B could be hypnotised immediately into BII who could be dissociated into BIa and BIVa (after Prince, 1906, p. 520).

WHO IS CONSCIOUS NOW?

• SECTION SIX: SELF AND OTHER

<i>Miss Beauchamp</i>	<i>Chris/Sally</i>	<i>Miss X</i>	Some cases could not have been created by therapy, such as that of Ansel Bourne, who, as far as we know, never had any therapy. In any case, if even a few of these fascinating cases really happened as described, they should tell us something very interesting about the relationship between self, memory, and consciousness. But what?
I	You	-	
You	I	She	The distinction between ego and bundle theories may be helpful here. Prince was clearly an ego theorist, for he believed not only in the existence of the 'real Miss Beauchamp' but in several other different selves who were distinct consciousnesses with separate wills.
She	You	I	

FIGURE 16.4 • The power of pronouns to create selves. According to Harré and Gillett (1994), three distinct systems of pronouns were used in Miss Beauchamp's speech. This means that one body housed three distinct selves, not because Miss Beauchamp had three selves inside her, but because three selves were discursively produced.

So his was a kind of multiple-ego theory: an in-between variant between the classic ego and bundle notions. Hodgson and Myers had similar beliefs, and, like many of their contemporaries, their ideas were rooted in spiritualist notions of mediumship, possession, and the idea of human personality as an entity that might survive bodily death (remember that they used the term 'personality' to describe a conscious entity, rather than a set of character traits). William James thought that cases like this, along with other hypnotic phenomena (Chapter 13), provided proof of a secondary consciousness or 'under self', co-existing with the primary consciousness. Indeed, he believed that 'The same brain may subserve many conscious selves, either alternate or coexisting' (1890, i, p. 401). As we will see later in the chapter, James thinks we have to acknowledge that our selves are bundles, but also wants to allow for a persistent core of sameness, so his view is somewhere in the middle too.

'The same brain may subserve many conscious selves'

(James, 1890, i, p. 401)

A more clear-cut bundle-theory interpretation of Miss Beauchamp's experiences comes from discursive psychology, a field built on the principle 'that the mind of any human being is constituted by the discourses that they are involved in' (Harré and Gillett, 1994, p. 104). Within this framework, the sense of self is a product of the use of the first-person pronouns 'I', 'me', and 'mine'. Philosopher and pioneer of discursive psychology Rom Harré and philosopher and neurosurgeon Grant Gillett use the case of Miss Beauchamp to illustrate 'the difference between the old idea of the self as something inside a person and the new idea of the self as a continuous production' (p. 110). Analysing conversations between Dr Prince and his patient, they argue that Prince made sense of Miss Beauchamp's utterances in terms of three independent pronoun systems. While BI spoke of herself as 'I' and Sally as 'you' or 'she', Sally referred to herself as 'I', and so on. This produces three distinct selves in the sense that each 'I' indexes the same body, but a different continuous sequence of events and morally responsible agent. Taking the 'discursive turn', there is no more to the selves than that. As Harré and Gillett put it, 'There are not three little egos inside Miss Beauchamp, each speaking up through her mouth. The speaking parts are all there is to it. They are the phenomenon, and these speaking parts are the selves' (p. 110).

This kind of theory has the advantage of not having to rely on mysterious entities called selves, but runs the risk of failing to say anything about consciousness. If the words are all there is, why do we have this compelling sense of a continuous and unitary self who is the subject of experiences? Or maybe words are plenty powerful enough to create this sense, and the problem is simply that we underestimate the power of language to effortlessly embed itself in every experience, even to the point of creating an experiencer.

So far we have divided theories of self crudely into two categories: ego theories, which entail some kind of continuing entity, and bundle theories, which do not. In their extreme versions, neither deals adequately with unusual cases like multiple personality, nor with ordinary self-awareness. On the one hand, extreme ego theories entail mysterious untestable entities. On the other, simplistic bundle theories do not explain why we feel as though we are a continuing entity. In this chapter we will consider some theories that try to avoid these shortcomings.

First, we must be as clear as possible about the difference between the two types of theory. For a bundle theorist, it makes no sense to ask how many selves may be inhabiting one body, or which self is the real one. We encountered this sort of question, and a sceptical objection to it, in our discussion of split brains in Chapter 6. Whereas most scientists try to answer the question ‘how many selves in a split-brain patient?’, for those who believe ego theory to be false, this is a nonsensical question. Derek Parfit (1987) asks us to imagine an experiment in which one hemisphere sees a red screen, and the other sees blue. When asked how many colours there are, both hands write ‘Only one’, but when asked to say which colour, one writes ‘blue’ and the other ‘red’. Now, assuming that this imaginary experiment worked as Parfit said it would, are there two streams of consciousness? Are there two conscious selves? Parfit concludes that there are indeed two separate streams of consciousness, one stream seeing red and the other seeing blue, but there are not two conscious persons who do the seeing. Why? Because only an ego theorist can count the number of persons involved. For a bundle theorist, there is no such thing as a continuous self who experiences the stream. So whether we consider split brains or whole brains, ‘the number of persons involved is none’ (1987, p. 20).

It might seem obvious that materialist scientists should agree with Parfit, accept Hume’s denial, and be bundle theorists. After all, if the brain consists of millions of interconnected neurons whose activity gives rise to behaviours, memories, and perceptions, then there is no need for an experiencing self as well. Yet, as we have seen, some scientists still try to count the number of selves in a split-brain patient, or ask whether multiple personalities are really separate selves.

The situation may be rather like that with the Cartesian theatre. While it is easy, intellectually, to deny the existence of a persisting experiencing self, it is another matter to accept all the consequences of such a view. Some classic philosophers’ thought experiments can bring these consequences to life. Remember that the point of thought experiments is not that they could be carried out, but that we use them to clarify our thinking, and to do that we must follow the rules exactly (Chapter 2).

‘the greatest scientific and philosophical riddle of all – the nature of the self’

(Ramachandran and Blakeslee, 1998, p. 255)

THOUGHT EXPERIMENTS WITH THE SELF

Imagine that in the middle of the night, without leaving any traces or doing any harm, a mad Martian scientist comes into your room, removes your brain, and inserts it into your friend John’s body, and then his brain into yours (impossible of course, but this is a thought experiment). In the morning you stir, your dreams recede, and you wake into full consciousness. But who has woken up? Have ‘you’ woken up in John’s body? Will you scream and protest, and hope you are only dreaming that you are in an unfamiliar room and have hairy legs and a bushy




FIGURE 16.5 • Reprinted from S. Law (2000), *The philosophy files* (London: Orion), p. 66, with permission.

beard? What, if anything, would John's body feel about the switch? Would it reject 'you', or welcome 'you'? Is there anything it's like to be John's body now, or to be 'your' body with John's brain in it?

If you think that both you and John will wake up in the 'wrong' body, then presumably you think that the conscious self depends on the brain and not the rest of the body. So in another popular thought experiment, the Martians scan the brains and then swap only the patterns of neural

information. This time, all your memories and personality traits are swapped over, but the brains stay in place. Now who is it who experiences the feel of the hairy legs and the beard? You or John? Is the experiencing self tied to the body, the brain, the memories, or what?

Ego and bundle theorists differ fundamentally in their responses to such questions. The ego theorist might say 'of course it will be me' (or 'of course it will be John') because the self must be associated with something, whether it is the body, the brain, personal memories, personality traits and preferences, or some combination. In other words, there has to be a right answer to the question 'who has woken up?' Ego theorists may try to find that answer by investigating the relationships between the conscious self and memory, personality, attention, or other brain functions, or between the brain and the rest of the body and the environment.

For the bundle theorist, this is all a waste of time because none of us is a continuous experiencing self. Yes, the person in the bed might scream and shout and be very unhappy and confused, but if you ask 'is it really me?', then you reveal your own confusion. There can be no answer to this question because there is no such thing as the 'real me'.

Are you an ego theorist or a bundle theorist? If you are not sure, this next thought experiment may help you find out.

Imagine that you are offered a free return trip, anywhere you want to go, in a teletransporter (very much like the *Star Trek* transporter). All you have to do is step inside a special cubicle and press the 'Go' button, whereupon every cell of your body is scanned and the resulting information stored (though your body is destroyed in the process). The information is then sent, at the speed of light, to your chosen destination and used to reconstruct an exact replica of you. Although this science-fiction idea is meant only as a thought experiment, some people believe that something like this may one day be possible (Kurzweil, 1999). We will return to this and other possible futures for our selves at the end of the chapter.

Since your replica's body and brain are in exactly the same state as yours were when scanned, the replica will seem to remember living your life up to the moment when you pressed the button. It will behave just like you, look like you, have your personality and foibles, and in every other way be just like you. The only difference is that this psychological continuity will not have its normal cause, the

continued existence of your body or your physical or social environment, but will depend on the information having been transmitted through space.

The question is – will you go?

Many people are happy to go. They reason that if their whole body is completely replicated, they won't notice the difference: they will feel just the same as before, and indeed will be just the same as before. Others refuse to go. Their reasons may not be as rational but may be more forcefully felt. 'This journey is not travelling but dying', they may say; 'the person who appears in Ibiza is just a replica, not the real me. I don't want to die'. It may be some consolation that after they get back from their trip, the replica will be able to take over their life, see their friends, be part of their family, finish their projects and so on, but still it will not really be 'me'. They do not accept, as the bundle theorist must, that it is an empty question whether 'they' are about to live or die (Parfit, 1987).

Some further thought experiments delve deeper. Imagine that the old you fails to be destroyed. In the futuristic fantasy of British neurologist Paul Broks, he presses the button for a routine trip to Mars only to be informed later that something has gone wrong. His replica is fine, but he is still here, in contravention of the Proliferation of Persons Act. Rather than allowing two Pauls to live, the original, which should have been destroyed, may have to be killed. 'Even Bundle Theorists don't want to die', he says as he awaits his fate (2003, p. 223).

Late in the first century, Plutarch imagined a ship being restored by replacing all of it, plank by plank. When does the Ship of Theseus stop being the same ship? Related questions are raised by the high-tech teletransporter. Imagine that only a few cells are replaced, or any proportion of them you like. Is there now some critical percentage beyond which you die and a viable replica is created in your place? If 50% are replaced, what would you conclude? Would the person who wakes up be half you and half replica? This conclusion seems ludicrous, but still you may be tempted to say that there must be an answer – the resulting person must really be either you or someone else. If that is how you think, then you are an ego theorist.

With this in mind, we may now explore a few theories of self. The examples given here in no way cover all possible approaches, but we have chosen those



ACTIVITY 16.1

The teletransporter

Imagine you want to go to the beautiful city of Cape Town for a holiday. You are offered a simple, free, almost instantaneous, and 100% safe way of getting there and back. All you have to do is step inside the box, press the button, and . . .

The box is, of course, Parfit's teletransporter. In making the journey, every cell of your body and brain will be scanned and destroyed, and then replicated exactly as they were before, but in Cape Town. Would you press the button?

To create a memorable exercise and encourage people to think more deeply, use a few chairs or tables to make the box and provide a colourful 'Go' button for a volunteer inside to press. Ask someone to stand by the button and say whether they would press it or not. What does everyone else think? Would they say 'Yes' or 'No'? Do not allow any 'Don't knows' (if people do not want to answer publicly, they could write down their answer). Do not allow quibbles over safety or any other details – this is a thought experiment and it specifies that the box is 100% safe and reliable. If anyone won't press the button, this has to be for some other reason than that it might go wrong.

Now ask for a volunteer who said 'Yes' and ask them to explain why. Others can then ask further questions to work out, for example, why this person is not worried about having their body completely destroyed. Do the same with a 'No' volunteer. Bear in mind that people's reasons for not going may involve their deepest beliefs about their soul, spirit, God, or life after death. It is helpful to remember this even while pushing people hard to explain what they mean.

After the discussion, find out how many people have changed their minds. In a course on consciousness, it is instructive to ask this same question after a few weeks or months of study, and for this purpose it is helpful for everyone to keep a record of their answers. They may change.



FIGURE 16.6 • Where's the University?

which seem to bear especially on the relationship between self and consciousness. In each case we can consider, first, whether the theory is an ego or bundle theory; second, how it accounts for the experience of *seeming* to be a unified and continuous self; and, third, whether it helps us to understand the nature of consciousness.

THOUGHT ITSELF IS THE THINKER

William James is the obvious starting point, for he wrote extensively about both self and consciousness, and his ideas are still widely respected today. James built his theory first and foremost on the way it *seems*. Central to the concept of personal identity, he said, is the *feeling* of unity and continuity of oneself; one's own thoughts have a warmth and intimacy about them that distinguishes them from others' thoughts – although where exactly to draw the boundary between one's own and others' thoughts is not always clear.

The universal conscious fact is not 'feelings and thoughts exist,' but 'I think' and 'I feel'. No psychology, at any rate, can question the *existence* of personal selves. The worst a psychology can do is so to interpret the nature of these selves as to rob them of their worth.

(1890, i, p. 226)

James divides the self into two ever-present elements: the 'me' and the 'I'. The 'me', the empirical self or objective person, includes three aspects: the material self (including body and possessions), the social self (including how we behave with and are seen by others), and the spiritual self (this may seem an odd name to us today, but it includes mental dispositions and abilities, religious aspirations, and moral principles). This last part of the empirical self, or 'me', includes subjective experience. Within the stream of consciousness, says James, there seems to be a special portion that welcomes or rejects the

SELVES, CLUBS, AND UNIVERSITIES

The theory that the self is just a bundle of sensations, or a stream of words, or a collection of events happening to no one, is not easy either to understand or to accept. To make the task easier, we can think about clubs or universities.

Suppose that the Bristol gardening club thrives for many years and then, for lack of interest, folds. The few remaining members put away the books, tools, and other club possessions, and move on to something else. A few years later a new gardening enthusiast starts the club up again. She retrieves the books, but re-designs the stationery. She attracts a few of the old members and lots of new ones, too. Now, is this the same club or a different one? If you think there must be a right answer, then you do not understand the nature of clubs. According to bundle theories, the self is a bit like this.



CONCEPT 16.2

rest, an 'active element' that receives the sensations and perceptions of the stream of consciousness, and is the source of effort, will, and attention (Chapter 7). It is something like a junction at which sensory ideas terminate and from which motor ideas proceed. He could hardly have described the audience in the Cartesian Theatre better.

But strangely enough, for James this audience is still only part of the 'me', not the 'I'. The 'I' lies beyond all this: it is the subjective knowing thought, or pure ego, the self that I care about, the *felt* nucleus of my experience. This is 'the most puzzling puzzle with which psychology has to deal' (1890, i, p. 330). He describes the two main ways of dealing with it in a way that should by now seem thoroughly familiar to us:

Some would say that it is a simple active substance, the soul, of which they are thus conscious; others, that it is nothing but a fiction, the imaginary being denoted by the pronoun I; and between these extremes of opinion all sorts of intermediaries would be found.

(p. 298)

James criticises both. The 'soul theory', he says, explains nothing, guarantees nothing, and lacks any positive account of what the soul may be. He rejects Plato and Aristotle's substantialist views, Descartes's dualism, and Locke's associationist theory. As for Kant's theory, the transcendental ego is just a 'cheap and nasty' edition of the soul, he says, and inventing an ego does not explain the feeling of the unity of consciousness: 'the Egoists themselves, let them say what they will, believe in the bundle, and in their own system merely *tie it up*, with their special transcendental string, invented for that use alone' (p. 370). Perhaps he means that believing wholeheartedly in bundles is so frightening that the ego string gets invented to conceal them.

At the other extreme, says James, those who side with the Humeans, claiming that the stream of thought is all there is, run against the entire common sense of mankind which insists on a real 'owner', a spiritual entity of some kind, or a real proprietor to hold the selves together. This 'holding together', concludes James, is what needs explaining.

How does James escape from inventing a real proprietor or a special string of his own? His well-known adage is that '*thought is itself the thinker*, and psychology need not look beyond' (1890, i, p. 401). 'The phenomena are enough, the passing Thought itself is the only *verifiable* thinker, and its empirical connection with the brain-process is the ultimate known law' (p. 346). What he means is this. At any moment, there is a passing thought (he calls this special thought the Thought) that incessantly remembers previous thoughts and appropriates some of them to itself. In this way, what holds the thoughts together is not a separate spirit

Have you heard the old joke about Oxford University? An American visitor asks a student to show him the famous and ancient University of Oxford. The student takes him to the Bodleian Library and the Sheldonian Theatre, to Brasenose College, Christ Church, and Lady Margaret Hall, to the Department of Experimental Psychology and the grand Examination Schools; he shows him Magdalen Bridge and students punting on the Cherwell. At the end of his extensive tour, the visitor says, 'But where is the university?' (Ryle, 1949, p. 16).

Do clubs exist? Of course. Do collegiate universities exist? Of course. But neither is something more than, or additional to, the events, people, actions, buildings, or objects that make it up. Neither is an entity that can be found. According to bundle theory, the self is like this.

'thought is itself the thinker'

(James, 1890, i, p. 401)



FIGURE 16.7 • According to William James, the continuity of self is an illusion. We think we are a continuous owner of our thoughts, when in fact a passing series of owners inherits the collection of thoughts and its ownership from the previous one.

or ego, but only another thought of a special kind. This Thought identifies and owns some parts of the stream of consciousness while disowning others. It pulls together thoughts it finds 'warm' and calls them 'mine'. The next moment another Thought takes up the expiring Thought and appropriates it. It binds the individual past facts with each other and with itself. In this way, the passing Thought seems to be the Thinker. The unity we experience is not something separate from the Thoughts. Indeed, it does not exist until the Thought is there.

James uses the metaphor of a herd and herdsman. Common sense rules that there has to be a herdsman who holds the herd together. But for James there is no permanent herdsman, only a passing series of owners, each of which inherits not only the cattle but the title to their ownership. Each Thought

is born an owner and dies owned, transmitting whatever it realised as its self to the next owner. In this way is the apparent unity created.



PRACTICE 16.2

AM I THE SAME 'ME' AS A MOMENT AGO?

As many times as you can, every day, ask yourself the familiar question 'Am I conscious now?', and then keep watching. As 'now' slips away, and things change around you, try to keep steadily watching, and wondering who is watching. Is there continuity of self as you remain aware? Can you see what that continuity is like? Or is there none?

The question is '**Am I the same "me" as I was a moment ago?**' What is really required is not asking (or answering) the question in words, but looking directly into how it seems.

Is James then a bundle theorist? He rejects any substantial ego, so we might assume so. And presumably he ought to step happily into the teletransporter because when the replica stepped out at the other side, a new Thought would immediately appropriate the memories and warm thoughts sustained by the replicated brain and so induce the same sense of unity and continuity as before. Yet James placed his own theory somewhere between the extremes, and criticised Hume for allowing no thread of resemblance or core of sameness to tie together the stream of consciousness. For James, the task was to explain both the diversity and the unity of experience, and he felt he had accomplished this with his 'remembering and appropriating Thought incessantly renewed' (p. 363).

Does his theory help us understand the nature of consciousness? Up to a point – and James himself tells us where that point lies. In the end, he cannot explain how

the stream of thought accompanies the stream of cerebral activity, nor why, as he puts it, 'such finite human streams of thought are called into existence in such functional dependence upon brains' (p. 401). In other words, the great chasm still yawns.

'The worst a psychology can do is so to interpret [...] selves as to rob them of their worth'

(James, 1890, i, p. 226)

NEUROSCIENTIFIC MODELS OF SELF

Many neuroscientists deliberately avoid talking about the self and self-consciousness (e.g. Crick, 1994). Others discuss self-awareness as a sub-category of awareness in general, and some consider how the self-concept develops and how it can go wrong (as in amnesia or blindsight). Only a few attempt to explain why the self seems to be a continuous agent and a subject of experience. Their most common strategy is to equate the self with one particular brain process or functional area of the brain.

Ramachandran suggests that his experiments on filling-in (Chapter 6) mean 'we can begin to approach the greatest scientific and philosophical riddle of all – the nature of the self' (Ramachandran and Blakeslee, 1998, p. 255). Part of the motivation for these experiments was Dennett's insistence that filling-in would have to be done *for someone* (for some viewer or homunculus), and that since homunculi cannot exist, filling-in does not occur. As we saw, some kinds of filling-in do occur. But the argument is not entirely false, says Ramachandran. Filling-in occurs for *something* rather than *someone*, and that something is another brain process – an executive process (Ramachandran and Hirstein, 1997).

Ramachandran considers MacKay's executive process (Chapter 6), and control processes located in frontal or prefrontal areas, but argues instead for the limbic system. The processes that best match what the self is traditionally supposed to do are those that connect motivation and emotion to action selection, based on an incoming set of qualia. Filling-in can then be seen as a way of preparing qualia for interaction with limbic executive structures. So our conscious experiences are the input to this executive system.

Consciousness is 'the self in the act of knowing'

(Damasio, 1999, pp. 9, 168, et al.)

Rama concludes that a single unified self 'inhabiting' the brain is an illusion, but also that 'It is not difficult to see how such processes could give rise to the mythology of a self as an active presence in the brain – a "ghost in the machine"' (Ramachandran and Hirstein, 1997, p. 455). Yet his theory seems to try to account for a real rather than a mythical self, and does so with the unexplained idea that qualia are inputs to specific brain processes.

Antonio Damasio draws on his studies of brain damage and psychopathology to distinguish the proto-self, the core self, and the autobiographical self. The sense of self, he argues, has a preconscious biological precedent in the simplest organisms. This *proto-self* is a set of neural patterns which map the state of an organism moment by moment. More complex organisms have 'core consciousness', which is not dependent upon memory, reasoning, or language, and is associated with the core self, 'a transient entity, ceaselessly re-created for each and every object with which the brain interacts' (1999, p. 17).

Extended consciousness entails more complex levels of organisation. Possibly present in other species, it develops fully over our lifetimes, building on working

memory and autobiographical memory to give rise to our *autobiographical self*. Damasio is clear that this self is not any kind of separate entity but is the you that is born as the story of your life is told. As he puts it, 'You are the music while the music lasts'; 'the owner of the movie-in-the-brain emerges within the movie' (1999, pp. 191, 313). His theory entails not only a movie-in-the-brain (Chapter 5) but the idea that neural patterns 'are displayed in the appropriate areas of the brain stem, thalamus, and cerebral cortex' to generate the feelings (p. 73). Yet Damasio insists that there is no need for a homunculus to watch the display. It is 'watched' by other brain processes. Take the example of burning your hand on a hot plate. Your body can be in a state of pain because of the neural patterns and nociceptive signals, but for *you* to *experience* the pain requires something more. Specifically, 'a process that interrelates neural patterns of tissue damage with the neural patterns that stand for *you*, such that yet another neural pattern can arise – the neural pattern of you knowing, which is just another name for consciousness' (p. 73). But there is no explanation of how the display accounts for subjectivity or how 'neural patterns that stand for *you*' can be consciousness.

According to Bernard Baars's GWT, the self-system is part of the context hierarchy that influences what gets onto the stage. Indeed, it is the dominant and unifying 'deep context'. Baars uses James's distinction between the 'me' and the 'I', distinguishing the self-concept (including values and beliefs about oneself) from the more fundamental self-system (including self as observer and self as agent). This self-system is fundamental because 'Consciousness inherently needs to interact with a self-system, at least if its information is to be reportable and usable' (1988, p. 344). In this way, self and consciousness stand in the relationship of context to content: self 'is knowledge that provides the framework for all conscious experience [. . .], an overarching context for the flow of conscious events' (p. 327).

In a later publication, Baars draws on Gazzaniga's idea of the interpreter and concludes that 'full consciousness may not exist without the participation of [. . .] prefrontal self systems' (2005a, p. 50). But because consciousness integrates all the brain's functions, information cannot be accessed by the 'observing self' (the executive interpreters in parietal and prefrontal cortex) without consciousness (p. 47). So Baars seems to be saying that on the one hand, the self may provide the context within which consciousness can exist, but that on the other hand, the self may depend on consciousness to be able to play any functional role at all.

Applying his method of contrastive analysis, Baars considers experiences in which the sense of the self is disrupted or abnormal, including fugue and multiple personality, as well as depersonalisation, a fairly common syndrome in which people feel themselves to be unreal or mechanical or not themselves and experience distortions of their body image. All of these self-alien phenomena, Baars notes, are precipitated by events that disrupt the stable dominant context, as his model predicts, and are associated with loss of autobiographical memory, as one would also expect from a loss of stability in the context. Also, the disruption may happen quickly, but recovery is slow because it means rebuilding the whole context.

In dissociative conditions such as fugue and multiple personality, different selves alternate because different context hierarchies vie for access to the GW. This means access to the senses and to autobiographical memory, and is required for any reportable conscious experience. Because there is only one GW, this seems to preclude the possibility (accepted by James, Prince, and others) of simultaneous consciousnesses.

Self is 'knowledge that provides the context for all conscious experience'

(Baars, 1988, p. 327)

Is GWT an ego or bundle theory? In GWT, the self is not an illusion; it *is* the most enduring level of the dominant context hierarchy. As with Gazzaniga's interpreter and MacKay's executive (oddly, Baars never seems to mention the top level of MacKay's schema, the self-supervisory system), the self-system is physically instantiated, and presumably one could, as with ego theories, count how many there were in a given brain. However, Gazzaniga, MacKay, and Baars should presumably be happy to step into the teletransporter because the physical systems ought to be completely reconstructed by the machine.

The continuity of self is also real in GWT. Selves persist because self-systems do – though both can break down and reintegrate, with time, effort, and stress (1988, pp. 343–344). But the unstated problem here is how a mass of neurons with changing interconnections and ephemeral activity can *be* a continuous experiencing self. So this theory faces all the problems that remained for James, and the additional problem of explaining how physical continuity translates into experience. Perhaps we should conclude with Baars's own words: 'You are the perceiver, the actor and narrator of your experience, although precisely what that means is an ongoing question' (1997b, p. 142).

In his neuronal update of GWT, Stanislas Dehaene disagrees with those who, like Damasio, think there is a necessary link between consciousness and self-consciousness. For him, being conscious of some aspect of oneself is just another form of conscious access to the workspace. Instead of the information being about colour or sound, it is about one of the various mental representations of 'me' – my body, my behaviour, my feelings or thoughts. When I reflect upon myself, the observed and the observing 'I' are simply encoded within different brain systems (2014, pp. 24–25). But this does not really explain the sense of continuity that one brain system seems able to confer compared to others.

Whether these neuroscientific theories are ego or bundle theories depends on whether they propose a continuing neural basis to the self. Some recent research on individual differences in the functional organisation of the brain suggests there may be such a basis. Functional connectivity patterns (especially those in the fronto-parietal network) can be used as a 'fingerprint' to reliably identify individuals from a large group, regardless of whether their brains are resting or engaged in a specific task (Finn et al., 2015). This suggests that brain activity might provide a basis for the self's continuity, and its distinction from others. Research on functional connectivity as a predictor of attentional ability (Rosenberg et al., 2017) and on how general intelligence emerges from differences in network architecture (Barbey, 2017) adds further weight to the idea that significant aspects of what we think of as self can be understood as linked to various kinds of persistent neural structures.

LOOPS, TUNNELS, AND PEARLS ON A STRING

'I am a strange loop', proclaims mathematician and cognitive scientist Douglas Hofstadter (2007). 'I am a mirage that perceives itself'. Famous for his book *Gödel, Escher, Bach* (1979), Hofstadter delights in recursive, self-reflexive, or loopy mathematical structures. He recounts a childhood shopping trip with his parents to try



FIGURE 16.8 • Who am I?

'I am a strange loop'

(Hofstadter, 2007)

'Deconstructing the "I" holds about as much appeal for a typical adult as deconstructing Santa Claus would hold for a typical toddler'

(Hofstadter, 2007, p. 294)

out that new invention, the video camera. He pointed it at his father and saw his face on the screen, then he pointed it at himself, and then he was tempted to point the camera at the screen itself. But he was so nervous that he actually asked the shopkeeper's permission and was told not to do it! 'This suspicion of loops just runs in our human grain,' he says (2007, p. 36). When he got home and played with the camera, he discovered strangely complex emerging patterns – but no danger.

The brain is full of loops. Some are simple like video feedback, but others are self-referential, like the sentence 'I am the meaning of this sentence,' or paradoxical like Escher's

Drawing Hands, in which one hand is drawn as though it is drawing the cuff on the wrist of another hand which is drawing the cuff on the wrist of the first hand which is . . .

Strangeness, claims Hofstadter (1979), arises when a system seems to twist around and engulf itself. This happens in 'tangled hierarchies', where it is possible to keep climbing from level to level only to end up where one started: 'A strange loop is a paradoxical level-crossing feedback loop' (2007, pp. 101–102). Because the brain is a tangled hierarchy, with multiple levels of symbolic representations and no definite top or bottom, it is full of strange loops. But what does it mean to say that *I* am one?

It is important to be clear which level of description one is using. Looked at in one way, the brain is full of dancing symbols perceived by other symbols, and this, says Hofstadter, is what consciousness is: the brain's loopy self-descriptions amount to a conscious self, with a deeply twisted-back-on-itself quality at its core. At this level, the self is not an illusion but is represented as a real causal agent. But if you shift down in viewpoint, then all these symbols are just non-symbolic neural activity. Then 'the "I" disintegrates. It just poofs out of existence' (2007, p. 294). In this sense, the self is an illusion or myth, but a myth we cannot live without, because it is central to all our systems of belief about ourselves.

The theory of strange loops is a bundle theory. Symbols are constantly dancing in the brain with no truly persisting experiencer, even though similar self-referential loops may be constructed over and over again. Hofstadter talks about souls and how they become attached to their own bodies, but his 'soul' is an abstract structure within the brain, not a separate dualist entity. Its apparent continuity and unity are properties represented in the brain. And souls can be represented, with different degrees of fidelity, in many brains – the brains of everyone who knows me. Does this account for consciousness? Hofstadter claims that by accounting for the soul or the 'I' as a strange loop, he has also explained 'having a light on inside' or 'being conscious'. Consciousness simply is the dance of symbols. But if so, he does not explain why the dancing is or generates a what-it's-like.

German philosopher Thomas Metzinger also takes a representationalist view of self. One of nature's best inventions, he says, is an inner tool that he calls the phenomenal self-model (PSM). This is 'a distinct and coherent pattern of neural activity that allows you to integrate parts of the world into an inner image of yourself as a whole' (2009, p. 115). Because you have this self-model or self-representation,

you can experience your arms and legs as *your* arms and legs, experience certain cognitive processes in your brain as *your* thoughts, and experience certain events in the motor parts of your brain as *your* intentions and acts of will.

On Metzinger's 'self-model theory of subjectivity', self is the content of the PSM, and 'Consciousness is the *appearance of a world*' (2009, p. 15). This world seems to be a single and unified present reality, but what we see, hear, taste, and smell is limited by the nature of our senses, so that our model of reality is a low-dimensional projection of a much vaster physical reality. It is a virtual reality constructed by our brains. So our conscious experience of the world is not so much an image of reality as a tunnel through it.

Why then does it feel as though there is always someone in that virtual reality? How does it become an 'ego tunnel'? There are two reasons. First, our brain's world simulation includes an integrated inner image of ourselves, the PSM, that is anchored in bodily sensations and includes a point of view. Second, much of the self-model is 'transparent'. This may seem an odd term, but what Metzinger means is that we do not realise it is a model; instead, we take it for a direct window on reality. Just as we don't see the transparent lens when we look through a telescope, so we don't see the neurons firing when we look at the world around us. The self-model theory of subjectivity 'is that the conscious experience of being a self emerges because a large part of the PSM in your brain is transparent' (2009, p. 7).

PSM theory is a bundle theory in Parfit's terms. As Metzinger puts it, 'no such things as selves exist in the world: nobody ever *was* or *had* a self' (2003, p. 1). The impression that we are a persisting self is created by the PSM which models the self this way. As for consciousness, this '*appearance of a world*' is a very special phenomenon because it is part of the world and yet contains it at the same time. But what does the appearance of a world mean? If the cells in your LGN or early visual cortex construct a representation of the retinal image, does this mean a world appears? Or must that world appear *for* someone? PSM theory tries to explain how inwardness is created when reality appears within itself, and claims that it is this inwardness which accounts for subjectivity.

Like Hofstadter, Metzinger believes that explaining the nature of self explains subjectivity. This is similar to Graziano's attention schema theory, in which the self (including the body schema) is constructed as part of the attention schema, along with the world being attended to and the process of attention; the self is only a self model, but it is constructed by the same process that makes us conclude we are conscious.

After herdsman and tunnels, another striking metaphor is at the heart of British philosopher Galen Strawson's account of self: 'many mental selves exist, one at a time and one after another, like pearls on a string' (1997, p. 424). Strawson's pearls are particular patterns of neural activity, or states of activation, that come and go. He throws out the idea that either agency or personality is a necessary feature of the self and, most controversially, also denies that selves have long-term continuity over time. Each self may last a few seconds, or a much longer time, but then it disappears and a new one appears.

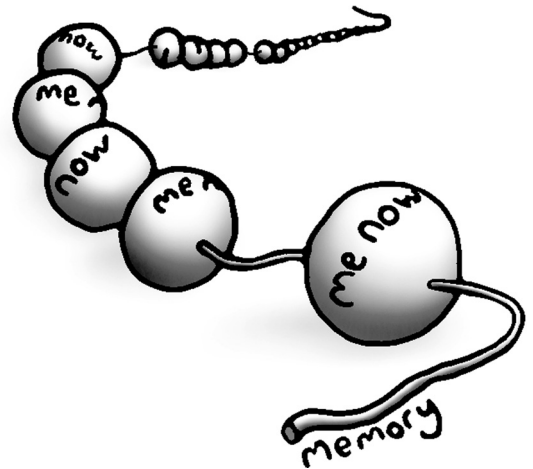


FIGURE 16.9 • According to the pearl view of self, many mental selves exist, one at a time and one after another, like pearls on a string.

'nobody ever was or had a self'

(Metzinger, 2003, p. 1)



FIGURE 16.10 • Perhaps there is no string that ties selves together and no requirement for only one to exist at a time. Selves may be more like raindrops, forming and disappearing again, sometimes lots at a time, sometimes only one. Being any one raindrop you would not know about the rest.

The stream of consciousness, '[I]like a bird's life, seems to be made of an alternation of flights and perchings'

(James, 1890, p. 243)

'the existence of the minimal subject is given with the existence of experience'

(Strawson, 2011, p. 254)

'the existence of a human being during a normal waking day involves the existence of many thin or minimal subjects'

(Strawson, 2011, p. 262)

Like James's, Strawson's theory depends on introspection, but he disagrees with James's description of 'the wonderful stream of our consciousness' which, '[I]like a bird's life, seems to be made of an alternation of flights and perchings' (James, 1890, i, p. 243). For him, even James's acknowledgement of discontinuity does not capture the radically disjunctive nature of experience. He prefers Hume's descriptions of consciousness as fluctuating, uncertain, and fleeting. There are gaps and fadings, disappearances and restartings, and he describes his own experience when alone and thinking as 'of *repeated returns into consciousness from a state of complete, if momentary, unconsciousness*' (Strawson, 1997, p. 422). It is as though consciousness is continually restarting.

People commonly accept that consciousness is 'gappy' or that thoughts switch and flip from one topic to another, says Strawson, but they still assume that the same self returns after a break. On the Pearl view, as in Buddhism, there is no such underlying continuity and no *persisting* mental self. The Pearl view is a radical version of bundle theory because of its complete rejection of any long-term continuity. Nevertheless, the pearl-self has unity at any given moment and in that sense is more than just an untied bundle of sensations and perceptions.

Does it account for the experienced unity and continuity of self? The pearl-self is similar to James's idea that each moment entails a new Thought. Its equivalent of the continuous appropriation of 'ownership', from one Thought to the next, is the suggestion that continuing contents help link experiences through time – the most reliable being the constant presence of our own body. Short-term memory also helps paste over the jumps and breaks. Thus, 'constancies and steadinesses in the *contents* of consciousness may seem like fundamental characteristics of the *operation* of one's consciousness, although they are not' (1997, p. 423).

Danish philosopher Dan Zahavi (2011) takes issue with one implication of Strawson's metaphor. If the identity of the self does not depend on temporal continuity, he says, we cannot call experiences pearls on a string, because two pearls are part of the same necklace only if they are in fact joined by an uninterrupted string. For Zahavi, there is no self that actively unites disparate bits of experience; it is not an extra element that has to be added to the stream of consciousness to unify it. Unity of experience in a single moment and through time is constituted by the fact that my experiences are mine, in the sense that an experience appears to me in a manner that is different from how they can appear to anyone else. And self in that sense exists whenever we are engaged or immersed in the world.

Perhaps because of confusions about continuity, Strawson (2011) drops the string-of-pearls metaphor, and expands on his earlier ideas by proposing that the self is a 'minimal self', or 'minimal subject': it is simply what remains when you strip away everything except experience, so if you agree that experience exists, you accept that the minimal self exists. A minimal self need not be self-conscious, and it may last for only a very short time (and so not be very ethically relevant, as he remarks in an aside). If ordinary daily experiences involve innumerable short and unnoticed gaps, then every day involves the existence of many minimal selves. Any given experience is by definition unified – it is a total experiential field – but (contrary to Zahavi's perspective) this unity does not imply anything about continuity. So maybe a raindrop would be a better metaphor.

I have very intense passions, and while I am in their grip my impetuosity is without equal: I know neither restraint, nor respect, nor fear, nor decorum; I am cynical, insolent, violent, bold: there is no shame that could stop me, nor danger that could frighten me: beyond the one object that occupies me, the universe is nothing to me. But all that lasts only a moment, and the moment that follows annihilates me.

(Jean-Jacques Rousseau, *Confessions [Les Confessions]*, 1782–1789)

SELVES AND BODIES, WORLDS, AND OTHERS

Some of the accounts outlined so far take some notice of the bodily and environmental context of the self; many do not. Zahavi's theory, influenced by the phenomenological tradition, emphasises the role of embodiment in structuring and constituting experience. A frequent collaborator of Zahavi's, Irish-American philosopher Shaun Gallagher, expands on this capacity as something given from the very outset:

The fact of embodiment is not something we need to reflect on to recognise. I do not need to reflectively ascertain that my body is mine, or that it is *my* body that is in pain or that is experiencing pleasure. In normal experience, this knowledge is already built into the structure of experience.

(2005, p. 29)

This raises again the question which came up in Zahavi's reflections on unity: if an important part of the concept of self is feeling that my experiences *belong to me*, we confront the question of what *mineness* is and feels like. Drawing on Buddhist ideas, the philosopher Miri Albahari (2006; and discussed in Zahavi, 2011) distinguishes between *perspectival ownership*, where the experience presents itself as distinctive to the subject of the experience, and *personal ownership*, a stronger sense of identifying oneself as the owner of an experience: thinking of it (whether reflectively or pre-reflectively) as being *mine* or apprehending it as part of *me*. Albahari thinks that having a sense of self requires not just perspectival ownership but personal ownership: not just being a point of view but drawing boundaries between what belongs to 'me' or other. But for her, crucially, having a sense of self is not the same as having a self. The sense of self exists but the self itself does not. Our experiences precede and create our sense of self rather than the other way round.

The idea that a self, or sense of self, is generated by facts about our experience (rather than experiences being possible because we have or are selves) is shared by social constructionist theories of self. These hold that selves are not born, but emerge in our interactions with other people. So instead of selves coming together in relationships, selves (or just the idea of selves) emerge from relationships.

'The other enters expressions of the self in their very formulation'

(Gergen, 2011, p. 647)

'Others think of me therefore I exist'

(Saunders, 2014, p. 93)

'consciousness does not really belong to man's individual existence but rather to his social or herd nature'

(Nietzsche, 1882/1974, trans. Kaufmann, §354)

'My body is an object all right, but my self jolly well is not!'

(Farrell 1996, p. 519)

Language is a crucial mediator of this emergence. The function of language is not to encode a set of neural representations; saying I am angry 'is more like a handshake or an embrace than a mirror of the interior' (Gergen, 2011, pp. 646–647). To say something is to perform an action within a relationship, and so 'private feelings' are better thought of as public actions: 'it is not that one has emotions, a thought, or a memory so much as one *does them*' (p. 647). And just as we cannot make ourselves understood if we use words we've just made up, so our actions do not make sense unless they draw on cultural traditions. Thus all our performances of self carry a history of relationships, and extend that history: 'The other enters expressions of the self in their very formulation' (p. 647). This is true even – or especially – for prisoners held in solitary confinement, who survive only by creating social worlds for themselves and by knowing that there are others who remember them. In conditions like this, where the volume is turned down on the everyday, it becomes clearer than ever that 'Others think of me therefore I exist' (Saunders, 2014, p. 93).

If we follow this path consistently, 'the other' stops seeming to be *outside*, but becomes part of the self. Selves don't first exist and then have intersubjectivity or sociality added to them: those qualities are just as intrinsic to them as embodiment. These qualities also connect directly to embodiment, as is clear in various important developmental activities in human children, such as imitation (learning by copying others' physical movements) and joint attention (attending to something along with someone else, like a mother and daughter reading a book together). In ways like these, we act with others, attend to others, and attend *with* others, and this way our shared experiences are part of who we are. And so the self is fundamentally 'dialogical' (Hermans, 2011).

Most of these embodied and extended theories are ego theories, in that there is continuity to their proposed selves. But they expand the boundaries of the ego so far that the self/other and self/world distinctions begin to dissolve, and it becomes hard to tell where I stop and the world begins. For Andy Clark (2008), the self extends beyond the boundaries of consciousness and beyond the skin, so that external resources like the information on my phone are a central part of my identity. Similarly for Alva Noë, the self is 'distributed' through the actions that connect my body with objects in the world: 'a person is not a self-contained module or autonomous whole'; a self isn't like a berry, but like the whole plant rooted in earth and tangled in brambles (2009, p. 69). This means the senses of unity and continuity no longer need explaining in their own right, since they are more closely tied to the unity and continuity that characterise the physical and social world in general. And if it doesn't seem this way to us – careful attention to the nature of experience may help it to.

CENTRES OF NARRATIVE GRAVITY

Some of the theories considered so far try to explain what the self is – a working part of the mental theatre, a special neural process, a strange loop, a mental model, or a consequence of embodiment or intersubjectivity. Others abandon the idea of the 'self itself' and try to account only for the feeling we have of having a self. Some hesitate between these two positions.

The last category we consider is those that put language at the centre, and here too we will see a tension between attempts to explain how and why we have *selves* (i.e. ego theories) or how and why we have mere *senses of selves* (i.e. bundle theories). What role is language meant to play in self-making? We have already considered its crucial social function. The other is its role as the medium through which narratives of self are created.

Narrative views of the self take many forms (Schechtman, 2011). Some say that our *sense of self* is narrative in structure, others that the *lives of selves* are. Some assume that selves must be agents, and that narrative is necessary for agency: a narrative context is what makes our actions meaningful and interpretable to ourselves and others. For example, a man's behaviour could be characterised with equal truth and appropriateness as digging, gardening, taking exercise, preparing for winter, or pleasing his wife (MacIntyre, 1985, p. 206). Which description the man chooses, for this and every other action in his life, determines who he is.

Such theories often leave it unclear whether and how we actually tell our self-narratives; if we are not aware of telling ourselves (or anyone else) a story, does that matter for a narrative theory? The question of how exactly narrative self-construction relates to language is also tricky, because it might result in denying consciousness, or perhaps just 'higher' forms of self-consciousness, to humans and other animals who have no language or only very basic language.

In most narrative theories, selves are the protagonists of the stories we spin – whether the story is of a horticulturalist or a good husband. Selves are sometimes also thought of as taking on more complex combinations of character, author, and critic in their own lives (Schechtman, 2011). Think of the selves you create via Facebook, Twitter, or email compared to those you create in face-to-face interactions with people. Not only are they all different narratives of yourself, mediated in different proportions by words, images, and bodily signals, but each new twist of each narrative affects the others, and affects the 'you' who creates the next one. The pressure to create ideal versions of yourself and the ways you impose commentary (sometimes ironic) on your own and others' idealised versions add further complexity to the feedback loops.

For Daniel Dennett, the self is a 'centre of narrative gravity'. But for him, unlike most other narrative theorists, this self is fictional in the strongest sense: there is no such thing as a self. When Dennett says that there is no Cartesian theatre, no show, and no audience, he really means there is no inner observing self. He claims that *'if you leave the Subject in your theory, you have not yet begun! A good theory of consciousness should make a conscious mind look like an abandoned factory'* (2005, p. 70).

If there really is no one in the factory, then Dennett must explain why we feel as though there is, and explaining how we come to believe falsehoods about consciousness is one of Dennett's favourite pastimes, as we have seen with zimboes, qualia, and vision. Do selves exist? Of course they do – and of course they don't! he says. There is obviously something to be explained, but not by invoking Ryle's 'ghost in the machine' or any mysterious entity controlling our bodies. So what kind of existence is it? For Dennett, self is a centre of gravity: it is invisible but real, a 'centre of narrative gravity'.

'the "Self of selves", when carefully examined, is found to consist mainly of [...] peculiar motions in the head or between the head and throat'

(James, 1890, i, p. 301)

'the trouble with brains, it seems, is that when you look in them, you discover that there's nobody home'

(Dennett, 1991, p. 29)

One problem, claims Dennett, is our tendency to think about selves as all or none, existent or non-existent. But just as we can be comfortable with fuzzy boundaries between species (cabbages and Brussels sprouts?), or between living and non-living (viruses?), so we should be with selves. They are biological products like spider's webs or bowerbird's bowers. They appeared gradually during evolution, and they are built gradually in each of our lives. Every individual *Homo sapiens* makes its own *self*, spinning a web out of words and deeds to build a protective string of narrative. Like spiders and bowerbirds, it doesn't have to know what it's doing; it just does it. The result is a web of discourses, without which an individual human being is as incomplete as a bird without feathers or a turtle without its shell.

But perhaps it is wrong to say that 'we' build the narrative. We humans are embedded in a world of words, a world of memes that are apt to take over, creating us as they go (Chapter 11). As Dennett puts it, 'Our tales are spun, but for the most part we don't spin them; they spin us. Our human consciousness, and our narrative selfhood, is their product, not their source' (1991, p. 418). This echoes the causal reversal we have seen proposed in several other theories – and this is where the 'centre of narrative gravity' comes in. When we speak, we speak *as if* the words come from a single source. They may be spoken by a single mouth, or written by a single hand, but there is no single centre in the brain (or mind, or anywhere else) from which they come. Yet we end up speaking as though there is. Who owns your car? You do. Who owns your clothes? You do. Who owns your body? You do. When we say 'This is my body', we do not mean the same as 'This body owns itself'. Thus, our language leads us into speaking and thinking *as if* there is someone inside: the audience in the Cartesian theatre, the 'central meander', or the inner agent. This self may be an abstraction, but, like the physicists' centre of gravity, it is a wonderfully simplifying and useful abstraction. This is why we have it.

For Dennett, multiple personality seems strange only because we falsely think that selves are all-or-none and must exist one to a body. Abandoning this idea allows us to accept fragmentary selves, partial narratives, and multiple selves that are just as real as the more common one-to-a-body type (Humphrey and Dennett, 1989). There might even be fewer than one self to a body, as in the case of the twins Greta and Freda Chaplin, who seemed to act as one, and speak together or in alternation.

Like Parfit, Dennett rejects the idea that in split-brain cases there must be some countable number of selves, but he goes further. 'So *what is it like* to be the right hemisphere self in a split brain patient?', or as Koch (2004) asked, 'How does it *feel* to be the mute hemisphere?' This, he says, is a most natural question, conjuring up a terrifying image of a self desperate to get out, but unable to speak. But this is a fantasy. The operation doesn't leave an organisation robust enough to support a separate centre of narrative gravity. The most it leaves is the capacity, under special laboratory conditions, to give split responses to particular predicaments, temporarily creating a second centre of narrative gravity. That this self could have gaps should come as no surprise. As Dennett has it, both self and consciousness may appear to be continuous but are in fact thoroughly gappy. They can lapse 'into nothingness as easily as a candle flame is snuffed, only to be rekindled at some later time' (1991, p. 423).

'Our tales are spun, but for the most part we don't spin them; they spin us.'

(Dennett, 1991, p. 418)

Dennett's is a bundle theory in which the string is a web of narratives. The sense of unity and continuity is an illusion abstracted from real words and deeds to the false idea of a single source. Does this help us understand consciousness? By denying that there is anything it is like to be an experiencing self, Dennett changes the problem completely. For him, the sense of self does not consist of an actual what-it's-like; it is nothing more than a centre of narrative gravity. Thus, language doesn't help create the experience, as it does for most other narrative theorists; instead, it creates the illusion of experience.

This is how he manages to explain consciousness – or alternatively, as some critics prefer to say, explain it away.



FIGURE 16.11 • Maybe a self can be snuffed out like a candle flame, and rekindled later. Maybe this is happening all the time even though we do not realise it (see Chapter 18).

FUTURE SELVES

initial downloads will be somewhat imprecise. [. . .] As our understanding of the mechanisms of the brain improves and our ability to accurately and noninvasively scan these features improves, reinstantiating (reinstalling) a person's brain should alter a person's mind no more than it changes from day to day.

(Kurzweil, 1999, p. 125)

For Kurzweil, and some other futurists, human selves will one day not be tied to the survival of human bodies: our immortality will be assured by technological progress. All we need to do is to increase the speed and accuracy of the scanning processes already available, copy the relevant aspects of a brain's organisation into a computer, and – hey presto – we live on. This dream is brought to life in the 2014 film *Transcendence*, where Johnny Depp plays an AI researcher who uploads his brain onto a quantum computer so that his consciousness can survive his body's death. As Kurzweil notes, we all change from day to day anyway, so a quick shift from bio- to silicon body should hardly be noticed.

Although such prospects have long been confined to thought experiments like those we considered earlier in the chapter, some people now think that it might really happen, and perhaps we should prepare ourselves. We may ask two questions. First, will the resulting creature be conscious? And second, will it be the same conscious person as before? Answers to the first question depend on whether you think there is something special about biology, or whether organisation alone is sufficient (as in functionalism). Answers to the second question depend on whether you are an ego or bundle theorist. If the opportunity ever comes, you may need to decide whether the operation really will make you immortal or not, but perhaps by then enough people will already have been copied, and will be telling you that it's fine and that they still feel just the same, for you not to care.

Kurzweil is, according to Rodney Brooks, one of those 'who have succumbed to the temptation of immortality in exchange for their intellectual souls' (2002, p. 205). According to Brooks, 'We will not download ourselves into machines; rather, those of us alive today, over the course of our lifetimes, will morph ourselves into

machines' (p. 212). To some extent this is already happening, with hip replacements, artificial skin, heart pace-makers, and cochlear implants. These electronic devices cannot yet match the sensitivity of a real human cochlear, or the number of connections to the brain, but they already enable profoundly deaf people to hear a good range of sounds, and even to enjoy music. Retinal implants are more difficult because of the number of neurons that join real retinas to their brains, but they are now available too: electrodes implanted in the retina detect light rays falling on to the retina, and convert them into electrical pulses which travel along the optic nerve to the brain. This allows blind people to read signs, tell the time on clocks, and distinguish red wine from white.

'The distinction between us and robots is going to disappear'

(Brooks, 2002, p. 236)

Replacements, or enhancements, for other body parts may be all metal and plastic, but they may alternatively be made from organic tissue, grown specially outside the body. 'Bioprinting' is an extension of 3D printing using plastic, human stem cells, water, and biocompatible material mixed to create living human tissue that can be matured into skin, liver, kidney, and other tissue types. Some severely disabled people can already control external devices by thinking, and some patients with locked-in syndrome are now able to communicate with the outside world. This is made possible by implanted electrodes that detect brain activity in motor cortex and use the signals to control wheelchairs, robots, or a computer mouse. 'The distinction between us and robots is going to disappear', says Brooks (2002, p. 236).

Imagine now a more exciting possibility: rather than a cochlear or retinal implant you can have an extra memory chip, an implanted mobile phone, or a direct brain link to the internet. Fanciful as these may seem at the moment, they are clearly not impossible, and would have implications for consciousness if they came about. So some speculation may be interesting.

Let's consider first the memory chip. Suppose that you have tiny devices implanted in your brain and can buy vast quantities of information to load into them. Since they have direct neural connections, the result is that your memory is vastly expanded. What would this feel like? It might, oddly enough, not feel odd at all.

Let me ask you a question. What is the capital of France? I presume that the answer just 'popped into your mind' (whatever that means) and that you have no idea where it came from or how 'your brain' found it. The situation with the memory chip would be just the same, only the world available to you would be greatly expanded.

Now add the implanted mobile phone, so that you can contact anyone at any time. With electrodes to detect your motor intentions, you could phone a friend any time by just thinking about them. And, finally, add permanent access to the web with search facilities and browsers all implanted in your head. With electrodes detecting your intentions, you would only have to think your question clearly enough and the answer would pop into your mind, just as the word 'Paris' did.

What would it be like to be such an enhanced person? Perhaps it would seem as though the whole of the web is as good as part of your own memory. Much of what you find on the web is junk and lies, but then ordinary memory is like that too. The skill of navigating through the vastness of cyberspace would only be an extension of the skills of using ordinary, fallible memory now. The odd thing is that everyone would have access to a lot of the same material.

An interesting question then arises. Who, or what, is conscious? Is it you, the web as a whole, the group of people using it, or what? According to GWTs, information becomes conscious when it is made globally available to the rest of the brain. In this speculative future, the whole of the web is globally available to everyone. Does that mean it would all be conscious? And if so, to whom or what? The notion of 'consciousness as global availability' seems to provide a curiously literal conclusion here.

If consciousness is unified by a (real or illusory) self, then nothing much changes by adding more memory, but once people are so intimately linked with each other, the whole concept of self seems under threat. What would make an item of information 'my' memory rather than yours? Perhaps having a physical body is still the anchor to which a sense of self adheres, but that too may be threatened.

SELVES IN CYBERSPACE

A teenage girl, pretending to be a boy, flirts in an internet chat room. Most of the other people she meets there are just that – people – but some are chatbots: programs designed to generate text and appear to be people. Virtual warriors inhabit millions of home computers, winning and losing battles in countless games, and acquiring personalities that are known the world over. Virtual actors live and die in films. A virtual television presenter stands in the studio, enthusiastically introducing a real, live human. A neural network translates a sentence from one language to another by generating its own third language, or 'interlingua'. Crawlers amble around the world wide web collecting information on behalf of search engines or communications companies. They are autonomous and go where they like. All of these entities depend on physical substrates for their existence, but none has a permanent physical home. Could they be conscious?

These few examples raise again the question of what kind of thing can be said to be conscious. We often say that a person is conscious, or wonder whether our dog is. Nagel asked 'What is it like to be a bat?', not 'what is it like to be a computation?', 'what is it like to be a bat's idea of a bat?', or 'what is it like to be a virtual self?' David Edelman was sceptical that anyone would ever seriously pose the question 'what is it like to be an octopus tentacle?' Several authors have argued that consciousness can arise only in physical objects that have boundaries and interests of their own, such as organisms and robots (Humphrey, 1992; Cotterill, 1998). Perhaps this is not true. Here we mean not free-floating psychic entities or astral bodies, but the possibility of conscious software agents that exist without being tied to one particular physical body. For example, they might be distributed across many machines. What, then, would give such entities any coherence, such that they could reasonably be said to be conscious selves?

According to meme theory (Chapter 11), memes tend to clump together into memplexes regardless of the substrate supporting them. So we should expect increasingly well-structured memplexes to form in cyberspace and compete with each other for survival. They would be purely informational entities with



FIGURE 16.12

increasingly sophisticated barriers letting some kinds of information in and rejecting other kinds. If they began using self-reference, then other memes could take advantage of this, elaborating their concepts of self. They would be much like selfplexes: the same things we create when we use language that refers to self.

We should expect a future in which increasing numbers of artificial personalities communicate routinely with us, answering the phone, dealing with our banking and shopping, and helping us find the information we want. They will probably be increasingly difficult to distinguish from what we now call real people, and as we saw in Chapter 12, we will respond to them as though they are.

To some people, 'the presence or absence of phenomenal consciousness can never be more than a matter of attribution' (Franklin, 2003, p. 64). Stan Franklin predicts that future software agents and robots will be so capable that people will simply assume they are conscious. Then 'The issue of machine consciousness will no longer be relevant' (p. 64).

When this happens and these beings claim to be as conscious as you are, will you believe them?

The idea of self links up with every other topic we have considered in this book. It is integral to why we think there is a problem of consciousness in the first place. It seems to provide a subject for the feeling of what-it's-like-to-be . . . It is the entity that declares it couldn't be under an illusion about something as dear to it as *its own experiences*. It is what feels as though it resides in your brain, or at least somewhere in your head, whether in a comfy seat in your own private theatre or just in the sense you have that there is unity, all the time, with you at its centre. It is 'you' who pays attention, decides to act, and disappears in unconsciousness. It is you for whom consciousness must have evolved, and you who is convinced that being an octopus is not like being you, and that being a machine is like being nothing at all, let alone like being like you. You are the one who chooses – or not – to expand your mind with drugs or meditation, or to lose yourself and briefly be someone else in a story or a film. It is you who seems to be only half-present in dreaming, and lost to view in mental illness.

Is it?

Do you feel any differently about any of this now, having read this far?

Does the you who might feel differently feel any less solid?

'The self posits itself, and by virtue of this mere self-assertion it exists'

(Fichte, 1794–1795/1982, trans. Heath and Lachs, p. 97)

READING

Broks, P. (2003). To be two or not to be. In P. Broks, *Into the silent land: Travels in neuropsychology* (pp. 204–225). London: Atlantic.

A routine teleportation to Mars fails to vaporise the original traveller on departure, contravening the Proliferation of Persons Act.

Dennett, D.C. (1991). The reality of selves. *Consciousness explained* (pp. 412–430). London: Little, Brown.

Selves as centres of narrative gravity.

Gallagher, S. (Ed.) (2011). *The Oxford handbook of the self*. Oxford: Oxford University Press.

Gallagher's introduction (pp. 1–28) summarises the contributors' arguments, including on whether the self is embodied (e.g. chapters by Bermudez, Cassam, Henry and Thompson, Legrand, Tsakiris) or socially or narratively constructed (e.g. Schechtman, Hermans, Gergen); and whether self is minimal (Strawson, Henry and Thompson, Zahavi), even less than minimal (Metzinger), or does not exist (Siderits). See especially Strawson (pp. 253–278) and Schechtman (394–416).

James, W. (1890). The consciousness of self. In W. James, *The principles of psychology* (i, pp. 291–401). London: MacMillan.


James's chapter on the self is long, but it is worth reading even a little of it to get a sense of his ideas on the Thought and the thinker. We especially recommend pp. 298–301, 329–342, and his own summary on pp. 400–401.

Metzinger, T. (2009). The empathic ego. In T. Metzinger, *The ego tunnel* (pp. 163–173). New York: Basic Books.

Includes a summary of the book's argument so far, and material on mirror neurons, the developmental self, and intersubjectivity.

Parfit, D. (1987). Divided minds and the nature of persons. In C. Blakemore and S. Greenfield (Eds), *Mindwaves* (pp. 19–26). Oxford: Blackwell.

Split brains, egos and bundles, and the power of our beliefs.



The view from within?

CHAPTER

SEVENTEEN

Introspective Observation is what we have to rely on first and foremost and always. The word introspection need hardly be defined – it means, of course, the looking into our own minds and reporting what we there discover.

(James, 1890, i, p. 185)

What do you discover when you look into your own mind? William James was confident: *'Every one agrees that we there discover states of consciousness'*, he said. But a hundred-odd years later we might be inclined to raise a few awkward questions. What does looking mean? Who is looking into what? Does the looking itself change what is seen? Is there value in looking without reporting? Does reporting destroy what we are trying to describe? Can everything be reported when some experiences are supposed to be ineffable? How reliable are our judgements about our states of consciousness? Are states of consciousness even the kind of thing that reliable judgements can be made about?

These are difficult questions. In the course of this book, we have found several reasons to reject the metaphor of vision turned 'inwards': the more we learn about how the brain and the rest of the body function in their physical and social environments, the less space there seems to be for any inner/outer split, or any interior space where consciousness is created. Nevertheless, we might agree that looking into our own minds is an essential part of studying consciousness. We cannot study consciousness in the abstract, because the what-it's-like-to-be is what we are trying to explain. So whether we follow the tradition of calling it introspection (from the Latin *spicere* and *intra*, 'look' and 'within'), or find some more neutral term for it, we cannot run away from the exercise.

We have already met many examples of people attending to their experience and reporting what they find. These include the methods of trained introspection developed by Wundt and Titchener, as well as James's descriptions of the 'flights and perchings' in the stream of consciousness, of getting up on a cold morning, and of religious experiences. Then there are various introspections on the experience of self, Csikszentmihalyi's studies of flow, and numerous adventures into altered states. Clearly, this personal approach has a role to play in the study of consciousness. But what sort of role?

The study of consciousness is sometimes divided into two fundamentally different approaches: the objective third-person approach and the subjective first-person approach. Between these two there is sometimes added another: the second-person, or inter-subjective, approach (Thompson, 2001). This is concerned with, among other topics, the development of empathy between people, the roles of mirror neurons, imitation, and joint attention in the relations between two people, and theories of intersubjectivity and how self is constructed through relationships with others.

There has been fierce argument over whether studying consciousness is fundamentally different from studying anything else, and whether it therefore requires a completely different approach from the rest of science. At the extremes, some people demand a complete revolution in science to take in the mysteries of consciousness, while others insist that we need no new approaches at all. The argument takes two forms which are often confused but are worth distinguishing. One concerns first-person versus third-person *science*; the other concerns first-person versus third-person *methods*.

There are at least three problems with the notion of a first-person *science*. First, although there are probably as many variations on scientific practice as there are people who call themselves scientists, all are part of a collective activity in which data are shared, ideas exchanged, theories argued over, and tests devised to find out which works better. The results are then published for everyone else to see, and to demolish or build upon further. Science, in this sense, is not something you can do on your own, suggesting that there can be no *privately* first-person science. But perhaps science then starts to look as much like second-person as third-person practice.

Second, objectivity is valued in science because of the dangers of personal bias obscuring the truth. So when one theory is easier or more comforting than another, the scientist is trained to set aside prior beliefs and maintain an open mind in the face of the evidence, suggesting that subjectivity might be damaging to science. There are good reasons, however, for treating the goal of scientific objectivity with some scepticism: perhaps we should be more honest with ourselves and admit that the attempt to subtract our subjectivity can never entirely succeed, and that trying to understand subjectivity better would be a worthwhile aim.

Third, as soon as inner explorations are described or spoken about, those descriptions become data for a shared scientific enterprise. In this sense, there can be no first-person *data* (Metzinger, 2003).

All these are arguments against a first-person *science* of consciousness, but none of them necessarily rules out a role for subjectivity, experiential work, or first-person *methods* in third-person science. For example, even on the strictest falsificationist



DO WE NEED A NEW KIND OF SCIENCE?

The table below is an attempt to lay out the arguments between those who believe that we need a fundamentally new kind of science and those who do not. Dennett calls them the A team and the B team, but this is only a shorthand. No one has signed up to these teams, and in reality there are far more than two positions to consider. So don't take the table too seriously – just use it as a way to remember the main issues at stake. You might like to fill in your own answers, too.

The last row leaves room for uncertainty. Clearly, the B team believes that first- and second-person methods are essential, but it is not equally clear what the A team thinks, or should think (do they consider these methods valuable, even if not essential?). For those who agree with A on all the other statements, the simplest response is to say that first- and second-person methods are inessential and so to ignore them. But another response is possible. Even if you believe that all data are third-person data and there are no 'experiences themselves', you may still think that private practices such as personal intellectual work, training in attention and concentration, or meditation and mindfulness may provide especially valuable third-person data. Perhaps these should not be called 'first-person methods', but the name seems appropriate, though the A team would not want them to be confused with a 'first-person science'.

	A	B	Your answer
We need a new kind of science to study consciousness	No	Yes	
First-person data are reducible to third-person data	Yes	No	
Third-person methods leave something out	No	Yes	

CONCEPT 17.1

theory of science, there is a role for experiential work and personal inspiration in the process of generating hypotheses. Such inspiration has often happened in science, and this is entirely valid as long as the fruits of the individual's work can be publicly tested. There has also been a long history of the public reporting of subjective impressions. None of these counts as first-person *science* because their data were publicly shared. But they might be counted as first-person *methods* to the extent that they involve systematic self-observation or self-exploration.

We can now see the difference between arguing for a first-person *science* of consciousness and arguing for first-person *methods* in a science of consciousness. If we argue only for first-person methods, we may then ask whether those methods need to be fundamentally different from the methods used in any other sciences, such as psychology, biology, or physics, or whether they are basically the same.

Arguably the distinction is less clear-cut than it seems, however. A scientist might think that she can start to meditate and introspect as a self-contained activity for generating hypotheses, and that her scientific practice will otherwise remain unchanged. But it is likely that the whole scientific process which surrounds the meditation will inevitably be changed, because meditating will change her views on what is worth testing, how the testing should be carried out, what counts as relevant evidence, and so on. From this perspective, it is meaningless to ask precisely how many new methods you can incorporate before you have a 'new science', because any change in method immediately changes the science. Nonetheless, the distinction still helps us assess more precisely what may or may not need to be different about a science of consciousness.

What appears to give the arguments a special twist when it comes to studying consciousness, as opposed to photosynthesis or black holes, is that the subjectivity is itself the phenomenon we are trying to explain. Here we meet a familiar argument. If there really are two separate worlds – the mental and the material, the inner and outer – then a science of consciousness is different from any other science and needs special methods for examining these non-material phenomena. On the other hand, if dualism is false and the inner and outer, mental and material worlds are one, then a science of consciousness need be no different from any other science.

If you think that a science of consciousness must be a fundamentally new kind of science, then you probably think that special first- and/or second-person methods are what is needed, and that enough of these together will constitute a suitably new science. If you think that a science of consciousness must be basically the same as any other science, then first- and second-person methods may still be relevant, but you must ask what role they can play, and whether they have anything special to contribute.

Either way, it is worth learning more about these methods. They include training our powers of attention and observation, developing our ethical and spiritual lives, actively exploring altered states of consciousness, and simply spending time thinking and questioning. All these are forms of personal work which may, or may not, contribute to the public process of coming to understand the nature of consciousness.

In this chapter, we shall first consider the furious debates that have raged over the role of first-person methods, and then consider some of those methods themselves.

	A	B	Your answer
Introspection observes the experiences themselves	No	Yes	
Mary learns something new when she sees red	No	Yes	
We must avoid the zombic hunch	Yes	No	
The distinction between first-, second-, and third-person perspectives is a false distinction	Yes	No	
First- and second-person methods have an essential role to play	No	Yes	

THE BATTLE OF THE AS AND BS

'I'm captain of the A team', proclaims Dennett, 'David Chalmers is captain of the B team', and so begins the battle over what Dennett calls 'The fantasy of first-person science' and Chalmers calls 'First-person methods in the science of consciousness'.

For Chalmers, the science of consciousness is different from all other sciences because it relates third-person data to first-person data. Third-person data include brain processes, behaviours, and what people say, while first-person data concern conscious experience itself. He takes it for granted that there are first-person data.

It's a manifest fact about our minds that there is something it is like to be us – that we have subjective experiences – and that these subjective experiences are quite different at different times. Our direct knowledge of subjective experiences stems from our first-person access to them. And subjective experiences are arguably the central data that we want a science of consciousness to explain.

(Chalmers, 1999)

At the moment we have excellent methods for collecting third-person data, says Chalmers, but we badly need better methods for collecting first-person data. The science of consciousness must hunt for broad connecting principles between first- and third-person data, such as certain experiences going along with certain brain processes or with certain kinds of information-processing. This is a hunt not for the

'the development of more sophisticated methodologies for investigating first-person data [...] is the greatest challenge now facing a science of consciousness'

(Chalmers, 1999)

'consciousness has a first-person or subjective ontology and so cannot be reduced to anything that has third-person or objective ontology'

(Searle, 1997, p. 212)

correlates of conscious versus unconscious, but for the correlates of different types of processing and experience. What he calls a 'fundamental theory of consciousness' would formulate simple and universal laws that explain these connections. Yet, argues Chalmers, data about conscious experience cannot be expressed wholly in terms of measures of brain processes and the like. In other words, first-person data are irreducible to third-person data (Varela and Shear, 1999).

Along with Chalmers, the B team includes Searle, Nagel, Levine, Pinker, and many others. Searle (1997) agrees with Chalmers about the irreducibility, although they disagree about much else (Chalmers, 1997). Searle puts it this way: 'consciousness has a first-person or subjective ontology and so cannot be reduced to anything that has third-person or objective ontology. If you try to reduce or eliminate one in favor of the other you leave something out' (Searle, 1997, p. 212).

Searle asks us to pinch our own forearms. Do it now and see what happens. According to Searle, two totally different kinds of thing happen. First, neuron firings begin at the receptors and end up in the brain, and second, a few hundred milliseconds after the pinch, we experience the *feeling* or quale of pain. These are the objective and subjective events respectively, and one *causes* the other. By 'subjective ontology', Searle means that 'conscious states only exist when experienced by a subject and they exist only from the first-person point of view of that subject' (1997, p. 120).

According to Searle, the difference is not just epistemic – that you can *know* about your pain in a way that nobody else can – it is ontological: pains and other qualia have a subjective or first-person *mode of existence*, while neuron firings have an objective or third-person mode of existence. Others have argued that consciousness is epistemically but not ontologically irreducible. For example, Metzinger explains that our conscious experience 'is truly an *individual* first-person perspective. Our phenomenal model of reality is an *individual* picture. Yet all the functional and representational facts constituting this unusual situation can be described objectively, and are open to scientific inquiry' (2003, p. 589). In other words, if we knew everything that was going on in the body and brain, we could identify the individual's perspective, and there would then be nothing more to discover. For Searle, on the other hand, there is not just a subjective point of view; there are irreducible subjective *facts*, and these are what a science of consciousness has to explain.

'Searle's proposed "first-person" alternative leads to self-contradiction and paradox at every turning', claims Dennett (1997, p. 118). On his A team, he lists the Churchlands, Andy Clark, Quine, Hofstadter, and many others. For them, studying consciousness does not mean studying special inner, private, ineffable qualia, but means studying what people say and do, for there is no other way of getting at the phenomena, and when we really understand all the third-person facts about brains and behaviour there will be nothing else left to explain.

'Searle's proposed "first-person" alternative leads to self-contradiction and paradox at every turning', claims Dennett (1997, p. 118). On his A team, he lists the Churchlands, Andy Clark, Quine, Hofstadter, and many others. For them, studying consciousness does not mean studying special inner, private, ineffable qualia, but means studying what people say and do, for there is no other way of getting at the phenomena, and when we really understand all the third-person facts about brains and behaviour there will be nothing else left to explain.




FIGURE 17.1 • According to Searle's 'subjective ontology', two completely different things happen when you pinch yourself. There are the objective effects on skin and neurons, and the irreducible subjective fact of feeling the pain.

Here we meet another classic argument: the incorrigibility of the first-person view. The B team argues that we have privileged access to our own experiential states, Searle's 'subjective facts'; that is, only we can observe them and we cannot be wrong about them. The A team contends that we have privileged access only to how it

seems to us. Dennett suspects that 'when we claim to be just using our powers of inner *observation*, we are always actually engaging in a sort of impromptu *theorizing*' (1991, p. 67). This view is a precursor of more recent illusionist approaches to consciousness (Chapter 3). We are always creating fictions about our experiences, and it is only these fictions that a science of consciousness must explain.

According to Dennett (2001b), the B teamers fall for the 'Zombic Hunch': the hunch that there could be a creature physically and behaviourally indistinguishable from you but 'all dark inside'. He says of Chalmers:

He insists that he just *knows* that the A team leaves out consciousness. It doesn't address what Chalmers calls the Hard Problem. How does he know? He says he just does. He has a gut intuition, something he has sometimes called 'direct experience'. I know the intuition well. I can feel it myself. [...] I feel it, but I don't credit it.

For Dennett, then, falling for the zombic hunch is like going on crediting the intuition that living things have some kind of extra spark to them, or that the sun goes round the earth. So he asks, 'do you want to join me in leaping over the Zombic Hunch, or do you want to stay put, transfixed by this intuition that won't budge?' (2001b). He is optimistic that some time in the next century people will look back on this era and marvel that we could not accept 'the obvious verdict about the Zombic Hunch: it is an illusion' (2005, p. 22), and chuckle over the 'fossil traces' of today's bafflement about consciousness. For these future thinkers, it may still *seem* as though mechanistic theories of consciousness leave something out, but they will accept that, like the sun rising, this is an illusion. As for the zombic hunch, 'If you are patient and open minded, it will pass' (p. 23).

'First-person science of consciousness [...] will remain a fantasy'

(Dennett, 2001b)



PRACTICE 17.1

IS THERE MORE IN MY PHENOMENAL CONSCIOUSNESS THAN I CAN ACCESS?

Here is a task relevant to the distinction between P-consciousness and A-consciousness (Block, 1995, Chapter 2): is there more in conscious experience than can be accessed? This looks like a question for first-person inquiry because only *you* know the answer; *you* must look into your own experience to see whether there is more there than you can convey to anyone else or even describe to yourself.

You might like to look out of the window at a complex scene, take it all in consciously, and then try to access parts of it, for example by describing to yourself the objects you see, or counting the number of trees or people in the scene. Do you get the sense that when you access some parts of your experience, others disappear or become unavailable?

This exercise may have some strange effects. Try to get used to doing it before you consider the more intellectual question: can this first-person exercise tell us anything useful for a science of consciousness?

• SECTION SIX: SELF AND OTHER

This distinction between the A and B teams is only Dennett's way of having fun with the major differences, and it skates over many subtler distinctions between ways of explaining consciousness (Davies, 2008), but it still gets to the heart of a major gulf. We can hear echoes of familiar arguments: those about qualia, zombies, conscious inessentialism, AI, Mary the colour scientist, and the function of consciousness, to mention just a few. They seem to lie at the heart of a distinction that will not – so far – go away.

Chalmers distinguishes three types of view about consciousness: A, B, and C:

Type-A views hold that consciousness supervenes logically on the physical

Type-B are also materialist but reject logical supervenience on the physical

Type-C deny both logical supervenience and materialism.

Type-A views include eliminativist, behaviourist, and reductive functionalist views; type-B include nonreductive versions of materialism which hold that consciousness cannot be reductively explained even though it is physical; type-C include various kinds of dualism, in which some sort of phenomenal properties are taken to be irreducible. For the A-type, zombies are inconceivable and Mary learns nothing about the world (though she may gain an ability) when she comes out of her black-and-white room; for B-type, zombies are conceivable but metaphysically impossible and Mary does learn something; for C-type, zombies are possible and Mary learns something about non-physical facts. For Chalmers, even though A and B are both materialist and C is not, the gap between A and B (is consciousness logically supervenient?) is far greater than that between B and C




FIGURE 17.2 • The A and the B teams fight it out in the playground.

(is physicalism true?). For him, property dualism is the only reasonable option: type-B views are popular but not very coherent, and type-A are quite simply on the wrong side of the Great Divide between those who take consciousness seriously and those who do not. But Chalmers acknowledges that, as Dennett says, in the end he falls back on intuitions.

Ultimately, argument can take us only so far in settling this issue. If someone insists that explaining access and reportability explains everything, that Mary discovers nothing about the world when she first has a first red experience, and that a functional isomorph differing in conscious experience is inconceivable, then I can only conclude that when it comes to experience we are on different planes. Perhaps our inner lives differ dramatically.

(Chalmers, 1996, p. 167)

'I can only conclude that when it comes to experience we are on different planes'

(Chalmers, 1996, p. 167)

So the two teams end up either crediting or rejecting their own intuitions, being sure that consciousness either does or does not need its own special explanation and, in both cases, refusing to budge. Their exchanges amount to 'that schoolyard dialectic: "You've left something out!" "No I haven't." "Yes you have." "No I haven't." "Yes you have." etc. etc.' (Raffman, 1995, p. 294).

We will return to these differences, and to Dennett's proposed alternative of heterophenomenology, but first we need to look at some traditional first-person methods.

PHENOMENOLOGY

The term 'phenomenology' is used in several different ways. Sometimes it refers to a person's experience (*their* phenomenology; how it is for them) or to experiences in themselves (*the* phenomenology; or what it's like), but here we are concerned with phenomenology as a method and a philosophy.

As a method, phenomenology also has two meanings. In the broad sense, it refers to any methods for the systematic investigation of phenomenal experience (Stevens, 2000). In the narrower sense, it refers specifically to the tradition based on Husserl's philosophy, and its later developments by Martin Heidegger, Maurice Merleau-Ponty, Jean-Paul Sartre, and others. Here we are concerned not so much with the philosophy, which is often obscure and difficult for outsiders to understand, but with the methods that Husserl advocated for getting to the 'experience itself' (Gallagher, 2007, 2012; Thompson and Zahavi, 2007; Gallagher and Zahavi, 2012).

Husserl argued that there can be no meaningful distinction between the external world and the internal world of experience, and emphasised the importance of lived experience over scientific abstractions. In order to explore this lived experience, one should suspend, or bracket, all one's preconceptions and prior beliefs, especially those about the nature of the external world and its relationship to experience; one should step back from the natural attitude of observing a world 'out there', and into the phenomenological attitude which investigates the very experiences that we have. It does not matter whether things really exist, physically or objectively (for example, whether the apple you are looking at is there in front of you, or whether you are dreaming or hallucinating it); that is a question

for the natural sciences. What is at issue is the *phenomenon* of the apple, as it is constituted in your experience. He called this bracketing process the epoché (from the Greek for ‘suspension’ – the use of the acute or the grave accent is inconsistent). By starting with this procedure, he claimed to be able to study experiences openly, directly, and without tracing them back to what they refer to in the world. In other words, to describe without theorising.

This method of suspending judgement has much in common with traditional methods of meditation and contemplative training, and also with shifts in awareness that can happen spontaneously. Astrophysicist Piet Hut (1999) likens it to an experience with his first camera. After intensively taking photographs in his familiar home town, he seemed to have landed in a different world and to be seeing things ‘in a new light’. Indeed, he seemed to see the world *as light*. Anyone who has learned to paint or draw will recognise this experience. The learning seems to be less about how to use the pen, ink, or paint, and more about new ways of seeing, or how to look at things directly without being distracted by how you think they ought to be. In the same way, the phenomenological ‘gesture of awareness’ is about seeing the world anew.

Husserl’s aim was what he called an eidetic reduction (*eidetic* from the Greek *eidos*, ‘form’): a way of finding the essential features, or invariants, of people’s experience. He wanted to get ‘back to the things themselves!’, to the way things are actually given in experience, claiming that by providing precise and systematic descriptions of experience, we can discover the structure of consciousness.

The phenomenologist is helped, says Shaun Gallagher, ‘by the realization that consciousness is intentional. This is the first thing that we come to understand through the phenomenological reduction’ (2007, p. 687). In other words, all experience is experience *of* something or *about* something. Husserl calls this the ‘noematic’ aspect of experience, and the ‘noema’ is the object as it is experienced, which is part of the structure of intentionality. Note, however, that there are many spiritual traditions which would reject this fundamental assertion, claiming to find ‘pure consciousness’, or awareness without any object or any intentionality (Chapter 18).

Husserl’s project ran into many difficulties, and his theories have been long and hotly debated. His essential method of epoché has not been widely adopted, nor led to a science of experience on an equal footing with the natural sciences, as he hoped. Nevertheless, it has been used in various contexts to explore emotional states, or to describe what it is like to undergo certain experiences and so discover their ‘essence’ (Stevens, 2000). The typical method involves several stages of analysing interviews or written accounts of experiences. First comes the epoché, then a summary or narrative digest, then significant themes are extracted to find the fundamental constituents of that kind of experience in general.

Arguably this use of phenomenology is not a first-person method at all but a third- or a second-person one. Although the original intention was to explore lived experience by seeing through preconceptions, the actual method used depends on analysing what other people say. In this sense, it is no different from many kinds of psychology which use questionnaires, interviews, role-playing, and the analysis of written texts. The original intention of throwing oneself into a new way of being in the world seems to have been lost.

‘consciousness is intentional. That is the first thing that we come to understand through the phenomenological reduction’

(Gallagher, 2007, p. 687)

‘Reading about the epoché typically leads a student to contemplate the concept of the epoché, rather than really performing the epoché’

(Hut, 1999, p. 242)

Perhaps this is not surprising, for it is hardly easy to undertake a personal transformation by throwing off one's preconceptions and going beyond conceptualisation back to the things themselves. It is much easier to talk about it. As Piet Hut notes, 'Reading about the epoché typically leads a student to contemplate the *concept* of the epoché, rather than really *performing* the epoché (a danger Husserl kept warning about)' (1999, p. 242). In other words, the first-person method slips all too easily away.

A related problem is that much of the language of phenomenology is incomprehensible to those not steeped in it, and difficult language can make people give up on a different field before they even begin, which is always a shame. Phenomenology sometimes gives the impression of relishing linguistic complexity for its own sake. For example, this is how the French philosopher Natalie Depraz explains her use of the phenomenological reduction as an embodied practice.

I am proposing to bring to light a renewed reductive method, whereby the spectator is given a specific embodiment, and where the operation inherent in the reductive gesture is taken up again through the logic of its own reflexivity. By thus aggravating the oxymoron of the practical and the theoretical, internal to the reduction in its Husserlian heritage, my point is that, in fact, reflection and incarnation, contemplation and action are not opposed until each begins to fertilize the other, thereby intensifying each other to the point of becoming virtually indistinguishable from each other.

(Depraz, 1999, p. 97)

Depraz is using the standard vocabulary of her field; we are not suggesting she is an especially bad writer. But the vocabulary she uses risks taking simple ideas and making them very hard to understand. Perhaps she means that when you look deeply into the distinction between subject and object, or between thought and action, the difference seems to disappear. If so, this is something found in many traditions, and it is not entirely clear how phenomenology helps.

NEUROPHENOMENOLOGY

Neurophenomenology is the name given by Chilean neuroscientist Francisco Varela to a 'quest to marry modern cognitive science and a *disciplined approach* to human experience' (1996, p. 330). He agrees with Searle that first-person



PROFILE 17.1

Francisco Varela (1946–2001)



Born in Chile, Francisco Varela studied biology before moving to the USA for a PhD on insect vision at Harvard, and later worked in France, Germany, and the United States. He said that he pursued one question all his life: why do emergent selves or virtual identities pop up all over the place, whether at the mind/body level, the cellular level, or the transorganism level? This question motivated his work on three topics: autopoiesis or self-organisation in living things, enactive cognition, and the immune system. Critics claim that his ideas, though fluently described, make no sense, and even friends described him as a revolutionary who threw out too much accepted science. His Buddhist meditation, as a student of Chögyam Trungpa Rinpoche, informed all his work on embodied cognition and consciousness. Uniquely both a phenomenologist and a working neuroscientist, he coined the term neurophenomenology. Reflecting on his liver transplant, he wrote vividly of the shifting sense of body and boundaries (Varela, 2001). Until his death he was Director of Research at the CNRS Laboratory of Cognitive Neurosciences and Brain Imaging in Paris.

'a quest to marry modern cognitive science and a disciplined approach to human experience'

(Varela, 1996, p. 330)

experience is not reducible to third-person descriptions, but proposes a new way of dealing with this irreducibility. Chalmers's hard problem cannot be solved, he says, by piecemeal studies of neural correlates of experience, but requires a strict method for rediscovering the primacy of lived experience. To get past piecemeal correlations and pure theory, we need systematic exploration 'of the only link between mind and consciousness that seems both obvious and natural: *the structure of human experience itself*' (p. 330). Anyone following this method must cultivate the skill of stabilising and deepening their capacity for attentive bracketing and intuition, and for describing what they find.

Varela describes the basic working hypothesis of neurophenomenology as that 'Phenomenological accounts of the structure of experience and their counterparts in cognitive science relate to each other through reciprocal constraints' (1996, p. 343). So, the findings of a disciplined first-person approach should be an integral part of the validation of neurobiological proposals. This is perhaps the kind of coming-together that philosopher Dan Lloyd imagines in his novel about a theory of consciousness: 'a transparent theory of consciousness, a Rosetta stone – you'd put in phenomenology at one end and get spiking neurons at the other' (2004, p. 31).

What does neurophenomenology mean in practice? Varela suggests that as techniques for brain imaging improve, 'we shall need subjects whose competence in making phenomenological discriminations and descriptions is accrued' (1996, p. 341). The basic idea is to gain more accurate descriptions of experiences in order to correlate them with measures of brain activity.

The practice of neurophenomenology has gradually been finding its way into neuroscientific experiments. A 2002 study by Antoine Lutz, Varela, and colleagues is often cited as one of the early examples of neurophenomenology in action. The idea is to take individual variation seriously, rather than simply averaging out everyone's results and pretending they are all the same. Participants were presented with a 3D illusion, and first-person reports about the participants' mental states were elicited after every trial. These were used to identify phenomenological clusters, and for each cluster the EEG imaging results were analysed separately. The neural patterns turned out to correlate with the degree of cognitive preparedness and immediate perception of the illusion as reported verbally by the participants. This suggests that variation which would otherwise have to be written off as 'noise' can be meaningfully interpreted by treating participants' first-person experiences as valuable data in their own right. Of course, if consciousness is what we are investigating, this should come as no surprise. But this and later work (e.g. Garrison et al., 2013; Petitmengin et al., 2013) helps show the concrete benefits of integrating 'first-person' accounts with neuroimaging data.

The experiment has been criticised, mainly for the lack of detail provided about the phenomenological side of their procedures. The authors state that participants were 'trained extensively with a well-known illusory depth perception task', and that they 'underwent the task until they found their own categories to describe the phenomenological context in which they performed it and the strategies they used to carry it out' (Lutz et al., 2002, p. 1586). But their reporting is

especially opaque on how the first-person behaviours were collected and clustered into categories. The authors do not say how often subjects

described their experience one way or another, what assumptions were made in encoding the data, how many experimenters encoded the data, how much the encoders agreed in their clustering of the data, and whether the encoders were blind to the hypothesis being tested.

(Piccinini, 2010, p. 104)

Given that the protocols for reporting the imaging side of the study are so much better established, these gaps are perhaps understandable. But they do compromise the aim of developing 'interpersonal standards of data-gathering' to apply to the exploration of subjectivity (Dennett, 2011, p. 32). More recent experiments have given detailed accounts of the instructions participants received and how the verbal data from their reports were analysed. One example is a study of 'effortless awareness' that related expert meditators' descriptions of their experiences to activation in the default mode network, specifically in the posterior cingulate cortex, which is active during self-related thinking (Garrison et al., 2013). Another example is a visual masking study (Albrecht and Mattler, 2012) which found correlations between three distinct categories of participants' performance and their reports of their perceptual experiences, though the analysis of participants' free reports was limited to very basic presence-or-absence ratings (do participants mention a motion percept, an afterimage percept, both, or neither).

It's still early days, but philosophers Evan Thompson and Dan Zahavi (2007) have argued for the value of collaborative research between phenomenology and neuroscience for such topics as self-consciousness, non-reflective self-awareness, temporality, intersubjectivity, and the importance of embodiment in experiencing the world.

Take temporality. The sense of time is potentially a rich area for study, because experienced time does not equate to neural time, and all sorts of anomalies arise when we try to pin down the 'time at which consciousness happens' (Chapter 6). Perhaps disciplined first-person study of experienced time might help.

According to Varela, this means exploring 'the structure of nowness as such' or what James called the 'specious present'. As James and others have described it, there is a three-fold structure in which the 'now' is bounded by the immediate past and immediate future. Husserl explored what he called internal time-consciousness. To hear a melody, see something moving, or see it as retaining identity over time, consciousness must be unified in some way through time. He introduced the twin ideas of *retention*, which intends the just-past, and *protention*, which intends the immediate future. So when we hear and understand a sentence, for example, we not only retain what has just gone but have some 'protention' of where the meaning of the sentence is going.

One of the ways in which Varela tried to bring phenomenology and neuroscience together was by relating the structure of time as discovered phenomenologically to the underlying self-organising neural assemblies. He explains that 'the fact that an assembly of coupled oscillators attains a transient synchrony and that it takes a certain time for doing so is the explicit correlate of the origin of nowness' (1999, p. 124). Varela describes this insight as a major gain of his approach. Yet some meditators have described a complete loss of any sense of 'now' resulting from deep exploration of koans such as 'when is this?' or 'are you here now?' Disciplined

'Consciousness is then, as it were, the hyphen which joins what has been to what will be, the bridge which spans the past and the future'

(Bergson, 1920, trans. Carr, p. 9)




FIGURE 17.3 • Varela devised this two-dimensional scheme for categorising theories of consciousness. Use it to try to position as many theories as you can. Don't turn the page until you've done it.

'I can't grasp a moment from which to say that what has gone before is past and what is to come next is future'

(Blackmore, 2011, p. 95)

'If a bridge is to be built between the neural and experiential levels, it should be done where the river is shallow, where descriptions of mental processes are fine-grained on both sides'

(Petitmengin and Lachaux, 2013, p. 1)

attention to experience can lead people to perceive change occurring but with no apparent distinction between past and future (Blackmore, 2011). This suggests that a search for the origins of nowness may be a search for an interesting kind of illusion.

Varela's ambitions have been taken further by a new kind of brain imaging called iGBM (intracranial gamma-band mapping), which measures gamma-band activity between 40 and 50 Hz as a general index of neural processing. It has much greater temporal resolution than fMRI, and is precise enough in a single trial that the signal-to-noise ratio doesn't need improving by averaging across multiple trials. These features make it better for exploring singular experiences of individuals. But they also demand more precision from those individuals if the available resolution is to be useful. Claire Petitmengin and Jean-Philippe Lachaux (2013) argue that to maximise our chances of integrating the neural and the experiential, we must attend to the smallest temporal unit of experience, changing participants' focus of attention from *what* (e.g. what they are listening to) to *how* (how the experience changes over time, how much effort is involved, what its effects are, etc.). For them, the study of 'microdynamics' provides access to early and usually invisible stages of our cognitive processes,

where the distinction between the sensorial modalities, and between the 'subject' and 'object' poles seems to be less rigid than in later stages. We hypothesize that these early stages give us a glimpse on the process of co-constitution of subject and object, knower and known that is called 'enaction'.

(p. 5)

To help us understand where neurophenomenology fits into a science of consciousness, Varela (1996) provides a simple diagram with four directions in which theories of consciousness can go. He positions the best-known thinkers on it, but excludes quantum theories and dualism, and restricts himself to "naturalistic approaches": those which 'provide a workable link to current research on cognitive science' (p. 332). In the north, Varela places functionalist theories, suggesting that they are the most popular in cognitive science, and that they all rely entirely on 'third-person' data and validation. Opposite them, in the south, are the mysterians who claim that the hard problem is insoluble. In the east are the reductionists, epitomised by the Churchlands, and Crick and Koch, who aim to reduce experience to neuroscience. Opposite them, to the west, comes phenomenology, with an area cordoned off for those who believe that a first-person account is essential, including Varela himself.

This diagram is helpful for thinking about the relationships between different theories, and it also puts a spotlight on the role of first-person approaches in a science

of consciousness. Varela implies that there is a real difference between theories which take first-person experience seriously and make it essential to their understanding of consciousness, and those which do not. But are they really so different?

Baars thinks not. ‘We already have a systematic study of human conscious experience, and it is called “psychology”’ (1999, p. 216). He suggests that if we look at what psychologists have been doing for more than a century, we find that they have always studied the things that people say about their experience. Yes we need phenomenology, in the broad sense, but we do not need to start from scratch.

Varela claims that only theories within his cordon make first-person accounts essential, but is this really so? To consider examples from each quadrant, Nagel surely takes the first-person view seriously in developing his idea of what it’s like to be a bat, even though he concludes that we can never know. Crick, for all his extreme reductionism, talks about such aspects of consciousness as pain and visual awareness, and bases his theory on people’s descriptions of what they see. And Dennett, even though he is accused of denying consciousness or explaining it away, begins by describing his own experience of sitting in his rocking chair watching leaves rippling in the sunshine, and tries to account for *‘the way the sunset looks to me now’* (1991, p. 5). It turns out not to be trivial to divide theories into those that take the first-person view seriously and those that do not.

A REFLEXIVE MODEL

Some people reject the distinction between first- and third-person methods altogether. Max Velmans (2000, 2009) points out that all sciences rely on the observations and experiences of scientists. Scientists can discover objective facts in the sense of acquiring knowledge that is validated intersubjectively, but there are no observations in science that are truly objective in the sense of being observer-free. He proposes a thought experiment in which the subject and observer in a psychology experiment change places.

Imagine a participant looking at a light, and an experimenter studying her responses and her brain activity. We might say that the subject is having private first-person *experiences* of the light, while the experimenter is making third-person *observations*. But all they have to do is to move their heads so that the participant observes the experimenter and the experimenter observes the light. In this swap, nothing has changed in the *phenomenology* of the light, yet the light



ACTIVITY 17.1

Positioning the theories

Varela has positioned some of the best-known theories of consciousness on a simple two-dimensional diagram. Before looking at where Varela himself places the theories, try to use his diagram to do this task yourself.

For a class exercise, give each student a copy of the empty diagram and ask them to place on it every theory of consciousness they can think of, or do the exercise together on the board. This is a useful revision exercise and a good way of drawing together ideas from the whole course. Point out that there are no right answers. Although Varela devised the scheme, he is not necessarily right about where each theory should go. When everyone has filled in as many theories as they can, show them Varela’s version.

How well do they agree? Every discrepancy can be used to discuss the theories and to test students’ understanding of them. In addition, you might like to criticise the scheme itself. For example, are there really theories of consciousness for which first-person accounts are not essential?

Can you come up with a better scheme? For instance, you might try to position theories according to their answer to one of the big questions: is there a hard problem or not, is there a difference between phenomenal and access consciousness, is studying the brain the best way to study consciousness, is consciousness an illusion, are some animals conscious and not others, will machines ever be conscious (or are they already), does consciousness have a function. . . ? What other candidates are there, and which are the most helpful?

‘We already have a systematic study of human conscious experience, and it is called “psychology”’

(Baars, 1999, p. 216)

SECTION SIX: SELF AND OTHER

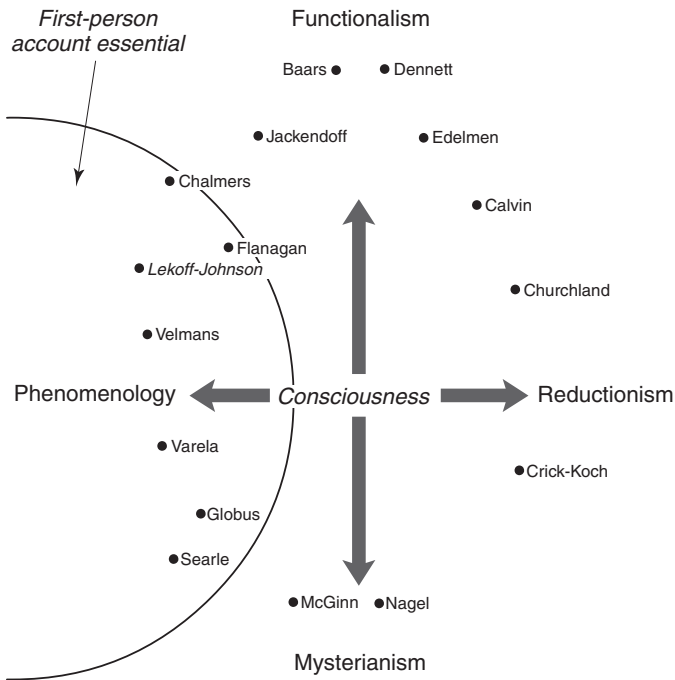


FIGURE 17.4 • Varela's 1996 categorisation of the major theories (after Varela, 1996, p. 332).

has supposedly gone from being a private experience to a public and objective stimulus. This, says Velmans, is absurd, and leaves us asking the fundamentally misguided question – is the light a subjective or an objective phenomenon?

Velmans thus rejects the dualism between subjective and objective phenomena, and proposes instead a 'reflexive model of consciousness'. He argues that our usual way of describing experiments misdescribes the phenomenology of perception, and hence misconstrues the problems facing a science of consciousness.

This reflexive model accepts conventional wisdom about the physical and neurophysiological causes of perception – for example, that there really is a physical stimulus in the room that our experience of it *represents*. But it gives a different account of the nature of the resulting experience. According to this nondualist view, when S attends to the light in a room

she does not have an experience of a light 'in her head or brain', with its attendant problems for science. She just sees a light in a room.

(1999, p. 301)

Explaining his theory of 'non-reductive, Reflexive Monism', Velmans argues that human minds and bodies are physical entities embedded in the universe of which they are a part, and capable of taking individual perspectives on the rest of the universe and themselves. As the universe evolves, it differentiates into parts which become conscious of themselves – hence the 'reflexive' aspect. Experience and matter are two sides of the same reality, as viewed from either a first-person or a third-person perspective. Causal links extend between the two, but neither can be reduced to the other: 'the contents of consciousness provide a view of the wider universe', but these 'conscious representations are not the thing-itself' (2009, p. 298).

Velmans claims that his model does away with many long-standing problems in the science of consciousness. He agrees that each of us lives in our own private, phenomenal world, and also that there are actual physical objects and events that people can agree about. But whether psychologists study mental or physical phenomena, they are doing so to establish repeatability and intersubjectivity. This, he claims, does away with the distinction between first-person and third-person methods. In both cases the motto should be: '*If you carry out these procedures you will observe or experience these results*' (1999, p. 300).

This motto is important. Think of the effects of drugs – 'if you take this drug you will experience these results' (Chapter 13) – or of practising mental disciplines: 'if you follow this procedure you will experience an OBE' (Chapter 15); 'if you meditate this way for many years you will gain this insight' (Chapters 7, 13, and 18).

'If you carry out these procedures you will observe or experience these results'

(Velmans, 1999, p. 300)

Yet reflexive monism faces serious problems. Velmans claims that it is 'nondualist' and calls it 'dual-aspect monism', and he concludes with the stirring idea that each of us is a small part observing the greater universe and so 'we participate in a reflexive process whereby the universe experiences itself' (2009, p. 298). Yet the theory rests entirely on the supposition that conscious experiences are 'private, subjective, and unique' and are constructions that represent the external 'things-themselves'. Thus, although reflexive monism is not a form of substance dualism, it seems to entail precisely the split that gives rise to the hard problem.

Velmans's ideas have nonetheless influenced other researchers interested in how to close the gap. Donald Price and James Barrell (2012) lament the fact that we do not yet have a true 'experiential neuroscience' (p. 19) and that human experience remains such a blind spot in the sciences – perhaps not accidentally. 'The maintenance of the blind spot of human phenomenal experience has not been a passive endeavor' (p. 26); rather, strenuous efforts have long been made by philosophers, psychologists, and neuroscientists to eliminate, reduce, and ignore experience.

They informally recreated Libet's famous (1985) wrist-flexing task, and found their colleagues reported a wide range of experiences while taking part, ranging from 'inner seeing' (an image of the hand moving), 'inner speech' ('I am going to move right NOW!'), emotional feeling (want to get this done!), and 'unsymbolised thinking' (the wordless equivalent of 'move real soon') – or else just 'I had no idea what was going on' or 'I was completely surprised by my hand moving'. Given such striking variety, they argue that Libet's methods and his interpretation of his results are far too simplistic, and do not represent 'the extended temporal phenomenology of choosing' (p. 286).

They suggest an alternative to Libet's experiment: get people to choose not when to flex their wrist, but how to cook a pizza, with either just a microwave or microwave plus conventional oven. Ask them not to respond randomly (which we never do in real life) but to deliberately choose one option and to notice what the choice feels like. How do you think this would change the experience of being a participant, and what we could learn as an experimenter? Like many others, they criticise Libet for using such a simplistic task, but their proposal departs far from the purpose of Libet's experiment, which was to measure the time difference between RP and W.

Price and Barrell offer another suggestion: imagine we have a complete mapping of neural activity correlated with a particular experience of pain, including the functional connections between activated and deactivated areas, and the interactions between the autonomic, somatomotor, and endocrine systems and the rest of the body. We also have a control condition without pain. The two are displayed on large screens and viewed 'objectively' 'by scientists who are disposed to leaving human experience out of the experiment altogether' (p. 26). What would these highly advanced scientists know? They would know that this is one pain and this is not-pain. But for deeper understanding, they would need fine-grained experiential maps of each, and the best way to get that would be for the scientists to be the participants. An account of how pain relates to neural activity requires observations of both: neither account can be observer-free, so why not, as Velmans suggested, have the same observer provide both? They could run the experiment on themselves and then use other people's accounts afterwards to confirm or disconfirm their direct observations.

'The maintenance of the blind spot of human phenomenal experience has not been a passive endeavor'

(Price and Barrell, 2012, pp. 25–26)

'almost direct observation of an association between brain activity and a specific type of experience by the same observer'

(Price and Barrell, 2012, p. 29)

The logical next step is to go beyond self-reflexive correlation to investigate causation through controlled feedback loops. This is already happening in studies that allow participants to change their own brain activity using visual feedback from real-time fMRI. In a nice example of how effective this can be, one study showed that by controlling activation levels in rostral anterior cingulate cortex, participants could change the intensity of pain caused by a noxious heat stimulus (deCharms et al., 2005). The pain did not change without the fMRI feedback, nor using different parts of the brain, nor with trick feedback from someone else's brain. So in this experiment we have 'almost direct observation of an association between brain activity and a specific type of experience by *the same observer*' (Price and Barrell, 2012, p. 29). Much more complex forms of neural interaction could be explored using this method, thus eliminating 'not human experience but *the perceived necessity of eliminating human experience* from science' (p. 29).

Real-time feedback loops in brain imaging are rapidly improving in precision. For example, the iGBM method outlined earlier can give feedback with much enhanced temporal resolution, letting us explore phenomena like mental imagery in greater depth (Petitmengin and Lachaux, 2013). More comprehensively reflexive experiments might even extend the feedback between brain and experience to cases where the gap between self and other is reduced or eliminated.

One example is the 'body swap illusion' (Petkova and Ehrsson, 2008). In this dramatic extension of the rubber-hand illusion (Chapter 4), the participant has a head-mounted display in which they see the input from a camera mounted on the experimenter's head. The two sit face-to-face, each holding a paintbrush in their right hand and using it to brush each other's left hand. This generates the illusion that the participant is brushing their own hand. You can also vary the procedure so the participant watches the experimenter's hand from the experimenter's perspective, or watches just their own hand (without the rest of the body) from the experimenter's perspective, or their own body (with or without the face visible) from the experimenter's perspective. What if we took this another step and gave the experimenter a camera, too? What if both participants could see their own and/or the other person's brain activity at the same time? What if both were highly trained neurophenomenologists? The prospects are exciting for reflexively integrating multiple perspectives into the neuroscientific study of consciousness and self.

These rules, the sign language and grammar of the Game, constitute a kind of highly evolved secret language composed of several sciences and arts, but especially mathematics and music (and/or musicology), and capable of expressing and establishing interrelationships between the contents and findings of nearly all disciplines. The Glass Bead Game is thus a play with the entire contents and values of our culture; it plays with them as, say, in the heyday of the arts a painter might have played with the colours on his palette. [. . .] Even if it ever happened that two players by chance should choose precisely the same small assortment of themes for the content of their Game, these two Games could, depending on

the way of thinking, character, mood, and virtuosity of the players, look and proceed completely differently.

(Hermann Hesse, *The Glass Bead Game* [Das Glasperlenspiel], 1943, our translation)

SECOND-PERSON NEUROSCIENCE

Paradigms like the body-swap illusion make clear that there is always a ‘second’ person in between the ‘first’ and the ‘third’. Second-person neuroscience is concerned with consciousness insofar as it asks how our conscious experiences relate to our attributions of consciousness to other people. Proponents suggest that some mainstream neuroscience has generated the results it has because its methods make them inevitable. For example, the science of ‘Theory of Mind’ is based on the idea that we engage in complicated inferencing and theorising about each other in order to bridge the gap between me and other people. By relying on people watching video recordings of other people and making judgements on what they see, for example, without ever testing their judgements in action or interaction, these studies found out only what they put in (Schilbach et al., 2013, pp. 394–395).

By contrast, advocates of the second-person paradigm tend to take what in Chapter 10 we called an ‘Interaction Theory’ approach to social cognition, which grew out of Gestalt theory and phenomenology. As one classic account puts it:

The quality of their actions imbues persons with living reality. When we say that a person is in pain, we see his body as feeling. We do not need to ‘impute’ consciousness to others if we directly perceive the qualities of consciousness in the qualities of action. Once we see an act that is skillful, clumsy, alert, or reckless, it is superfluous to go ‘behind’ it to its conscious substrate, for consciousness has revealed itself in the act.

(Asch, 1952, p. 158)

We immediately experience the other as a subject. The sense we have of a great gulf between myself and other people need not be thought of as an epistemological given, an inescapable limitation on the kind of science we do. Neuroscience, say Leonhard Schilbach and colleagues (2013), ‘should not content itself with a spectatorial view of social cognition’ (p. 443). Observing others and interacting with them are not the same thing, and when we ask questions about how humans typically interact, third-person observation doesn’t deserve to be scientifically privileged in the way it long has been. Neuroscience therefore needs new methods to encourage meaningful interaction between participants and between ‘participants’ and ‘experimenters’, including elements such as emotion and reward, nonverbal as well as verbal responses, the dynamics of real-time feedback, and more complex reconstructions of social encounters.

Better second-person methods of this kind might help ‘clos[e] the gap between the experiential and the neurobiological levels of description in the study of human consciousness’ (Olivares et al., 2015, p. 1). In second-person relationships, people

Neuroscience ‘should not content itself with a spectatorial view of social cognition’

(Schilbach et al., 2013, p. 443)

The view from within?

Chapter Seventeen

respond to us depending on our actions; in mere observer relationships, they do not (Longo and Tsakiris, 2013). This idea of action-based contingencies has become important in thinking about first-person experiences as embodied and enactive (Chapters 5, 6, and 8), and systematic feedback from the social environment may prove to be as powerful in shaping self-other consciousness as sensorimotor feedback seems to be for more 'private' forms of perceptual consciousness.



PRACTICE 17.2

SOLITUDE

Being alone is very rare, and the rarer it becomes the more potential there is for us to learn from it. Keep a whole day and night clear in your diary, and prepare in advance so that you will have no contact with anyone else: prepare all the food and drink you will need, tell people you will be out of contact, and switch off your phone and computer and all other electronic devices. The experience will be much more powerful if you avoid reading and writing too. If the only place you can be on your own is in a single room, do it there; if you can go alone into nature, even better. This task may seem daunting, and probably should: for most of us it is a big undertaking. But the difficulty is in direct proportion to the amount we stand to learn about ourselves from stepping outside our social selves for a day.

How does your sense of yourself, and of how time passes and how you relate to the world, change as the day wears on? How does the experience compare to your expectations? What is it like to be you, alone?

Blurring the boundaries between the different 'persons' becomes important here too, as in the reflexive model we explored earlier. In the 'enfacement illusion', for instance, seeing someone else's face being touched at the same time as your own changes your recognition of your own face and reduces the difference you see between theirs and yours. The effect also extends to things like seeing the other person as more attractive and tending to conform more to their behaviours (Paladino et al., 2010). And the size of the effect depends on how sensitive you are to your bodily states: less sensitive participants experience a stronger illusion (Tajadura-Jimenez et al., 2012). Together, findings like these suggest that the shared contingencies of social and multisensory interactions might explain how the self as subject and the self as object are tied together: 'how the "I" comes to be identified with "me", allowing this "me" to be represented as an object for others, as well as for one's self' (Longo and Tsakiris, 2013, p. 2).

If the problem is how to combine the benefits of first-person immediacy with third-person reliability, the second person could be part of the solution, for example if a trained interviewer helps people describe their experiences accurately. This mediator would not be distanced or neutral but would take an empathic stance allowing them to investigate an experience together with the participant. One proposal for avoiding bias is to make the mediator blind to the stimulus which the

participant is responding to – an idea known as the Double Blind Interview or DBI (Froese, Gould, and Barrett, 2011; Olivares et al., 2015).

Interviews conducted by the experimenters themselves were used in an experiment (Petitmengin et al., 2013) where participants were asked to choose which of a pair of portraits they preferred; on six out of fifteen trials, they were asked to explain their choice. Three out of the six times they were actually handed the non-chosen photo, but only 33% detected the deception, and a large majority gave explanations for why they had supposedly chosen this photo. But if instead of just being asked to explain their choice, participants had an elicitation interview, 80% *did* realise they had been misled.

In this experiment the interviewers were not blind, but all they did was assist in the act of remembering, rather than prompting any particular content, and the whole point of the interview was to ‘trigger the acts which enable the detection of truth’ (2013, p. 660), so blinding would have seemed perverse. The experimenters gathered rich descriptions (originally in French) which they then classified into different varieties of perceptual and decision-making experiences. These kinds of methods are still in their relative infancy as regards mainstream (neuro)cognitive research, but current progress is fast and creative.

HETEROPHENOMENOLOGY

Heterophenomenology (which might be translated as ‘the study of other people’s phenomena’) is an awkward name for our final method of studying consciousness. According to Dennett (1991, 2001b), it involves taking a giant theoretical leap, avoiding all tempting shortcuts, and following ‘the *neutral* path leading from objective physical science and its insistence on the third-person point of view, to a method of phenomenological description that can (in principle) do justice to the most private and ineffable subjective experiences’ (1991, p. 72).

Imagine you are an anthropologist, says Dennett, and you are studying a tribe of people who believe in a forest god called Feenoman and can tell you all about his appearance, habits, and abilities. You now have a choice. You can become a Feenomanist like them, and believe in their god and his powers, or you can study their religion with an agnostic attitude. If you take the latter path, you collect different descriptions, deal with discrepancies and disagreements, and compile as well as you can the definitive description of Feenoman. You can be a Feenomanologist.

This is possible because you are not treating Feenoman as a creature who might jump out from behind a tree and give you the *right* answers. Instead, you are treating him as an ‘intentional object’, a kind of fiction like Sherlock Holmes or Doctor Watson. In fiction, some things are true or false within the story, but others are neither. So, to use Dennett’s example, it is true that Holmes and Watson took the 11.10 to Aldershot one summer’s day, but it is neither true nor false that that day was a Wednesday because the author does not tell us. Similarly, there is no point in trying to find out whether Feenoman *really* has blue eyes: on this, and all other questions about Feenoman, the beliefs of the Feenomanists are authoritative, but only because Feenoman is being treated as their intentional object, i.e. a fiction. Their reports are authoritative only about how things *seem* to them.

AM I EXPERIENCING
MORE THAN I CAN
ACCESS?

The view from within?

Chapter Seventeen




FIGURE 17.5 • The feenomanologist collects his data from the Feenomanists.

WHAT IS IT LIKE TO BE YOU, ALONE?

This is the attitude that Dennett urges us to adopt in his method of heterophenomenology. It ‘neither challenges nor accepts as entirely true the assertions of subjects, but rather maintains a constructive and sympathetic neutrality, in the hopes of compiling a *definitive* description of the world according to the subjects’ (1991, p. 83). The people being studied may, like the Feenomanists, protest ‘But Feenoman is *real*’, ‘I really *am* having these qualia’, to which you, the heterophenomenologist, can only nod and reassure them that you believe they are sincere. This, says Dennett, is the price we have to pay for the neutrality that a science of consciousness demands. While heterophenomenologists accept people’s descriptions of how things *seem* to them, ‘we have to keep an open mind about whether our apparent subjects are liars, zombies,

or parrots dressed up in people suits’ (p. 83). An interesting comparison here is that while traditional phenomenology remains agnostic about the reality of the world, heterophenomenology remains agnostic about the reality of conscious experiences.

What does this mean in practice? Dennett describes the method in three steps. First, the data are collected. These might include brain scans, button presses, or people’s descriptions of mental images or emotions. Second, the data are interpreted. This step is unavoidable, and might include turning the brain scans into coloured pictures, relating the button presses to the stimuli presented, and turning the speech sounds into words that we write down and understand as descriptions of mental images. Third, and this is the crucial step, we adopt the intentional stance (Chapter 10). That is, we treat our informant as a rational agent who has beliefs, desires, and intentionality. We allow that she pressed the button because she wanted to tell us that she saw the green blob, and spoke those words because she was trying to describe her complex mental image or the powerful emotion she felt when you showed her that picture.

There may be inconsistencies that have to be investigated or ironed out, but in spite of these difficulties the method leads easily enough to the creation of a believable fiction: the subject’s heterophenomenological world.

This fictional world is populated with all the images, events, sounds, smells, hunches, presentiments, and feelings that the subject (apparently) sincerely believes to exist in his or her (or its) stream of consciousness. Maximally extended, it is a neutral portrayal of exactly *what it is like to be* that subject – in the subject’s own terms, given the best interpretation we can muster.

(Dennett, 1991, p. 98)

According to Dennett, this is the basic method that has always been used in the science of psychology, and he has not invented it but merely explained its rationale. Others claim that this kind of suspension of belief in the truth value of introspective reports is not as mainstream in the study of consciousness as Dennett makes out (e.g. van de Laar, 2008).

But isn't there something left out? Isn't there more in P-consciousness than we can access? Isn't heterophenomenology only studying what people say and leaving out *the experiences themselves*? Isn't it treating their inner world as a fiction when it *really exists*? Isn't it only treating people *as if* they have beliefs, desires, emotions, and mental images when they *really do*? Is Dennett – who is so obviously a real self with real consciousness – just fooling himself (and us) by pretending to have come up with a non-fictional account of why self and consciousness are fictions?

These questions get to the heart of the issue, so it is worth trying to be clear about how heterophenomenology responds. Heterophenomenology maintains neutrality on all these points. Just as in traditional phenomenology, theories are bracketed while the investigation proceeds. But whereas for phenomenology, the question being asked is 'why does this person experience X?'; for heterophenomenology, the question is 'why does this person say "I experience X"?' Heterophenomenology leaves open the question whether there is something more to be discovered, pending further investigation. One day we might discover a blue-eyed healer who was clearly the source of the Feenomanists' beliefs. Even if they resisted the identification, it might one day become clear, perhaps even to them, that the new guy was sufficiently like their old god to say we had discovered what they were really talking about – just as the goings-on in our brains may one day be understood well enough to say we could accept their identity with the phenomenology. Some diehard believers might still object that the *real* phenomenological items only accompanied the goings-on without being identical to them, but how much credence that claim should be given would be another matter.

While conducting their explorations, heterophenomenologists use the fiction of the heterophenomenological world much like a physicist might use the fiction of a centre of gravity, or the equator. They leave it open whether Feenoman really exists or not; whether *as-if* intentionality is different from *real intentionality* (Chapter 12). Dennett presumably thinks there is no difference, but heterophenomenology, as a method, is not committed either way.

What role remains for 'looking into our own minds'? Heterophenomenology has attracted much criticism from those who believe it is somehow opposed to the first person. But Dennett describes those who say they want a first- or second- rather than a third-person perspective as 'bickering over labels' (2007, p. 252). He says that 'heterophenomenology could just as well have been called – by me – *first-person science of consciousness* or the *second-person method of gathering data*' (p. 252). Indeed, 'Collaborating with other investigators on the study of your own consciousness (adopting, if you like, the "second-person point of view") is the way to take consciousness, as a phenomenon, as seriously as it can be taken' (Dennett, 2017, p. 351).

He chose the third-person label instead, he explains, to emphasise continuity with the objective standards of natural sciences, but 'the critiques are directed at the label, not the method' (2007, p. 252). He objects to what he calls 'lone wolf autophenomenology' (relying on oneself as the sole subject), and the ambition to found a 'single, unified *first-person science* of consciousness', which for him would amount to a '*solipsistic science*' (p. 264). But everything else that brings in first-person methods is good – and is, he claims, already heterophenomenology. Certainly one can adopt the heterophenomenological stance towards oneself in the reflexive way that Velmans and others advocate.

Heterophenomenology is 'the maximally open-minded intersubjective science of consciousness'

(Dennett, 2007, p. 264)

Heterophenomenologists need both more scepticism (about ourselves as well as our subjects) and more wonder (about what we are studying), says Dennett. Finding things out about someone's experience in a heterophenomenological way is different from having an ordinary conversation with someone because we have to maintain 'a deliberate bracketing of the issue of whether what they are saying is literally true, metaphorically true, true under-an-imposed-interpretation, or systematically-false-in-a-way-we-must-explain' (2007, p. 252). But we have to acknowledge that we can never be completely neutral, with *everything* bracketed off. We have to acknowledge that we cannot do anything without interpreting (for example, without adopting the intentional stance), and that whatever other methods might claim about not interpreting, this is impossible. Yet we also have to acknowledge that interpretation (contrary to what many humanities scholars believe) can be subject to rules and agreements.

Meanwhile, Dennett suggests, we should be more amazed at our ability to translate experience into report at all.

We tacitly take the unknown pathways between open eyeballs and speaking lips to be secure. Because we all can do it (those of us who are not blind) we don't scratch our heads in bafflement over how we can just open our eyes and then answer questions, with high reliability, about what is positioned in front of them in the light. Amazing! How does it work?

(Dennett, 2007, p. 255)

We have no more privileged access to this process than we do 'to the complicated processes that maintain the connectivity between our reporter's cell phone and ours,' says Dennett in a later book (2017, p. 349).

Explanation has to stop somewhere – and it tends to stop much sooner than we might think. This is clearer in the case of imagining rather than seeing, Dennett suggests: when we imagine something, we *know* we don't know exactly what we're experiencing or why or how to describe it. If you're still not convinced, try the example of an invented cognitive capacity. Imagine you can spread your toes and thereby come to have breathtakingly accurate convictions about what is happening in Chicago. And imagine not being curious about how this is possible. How do you do it? 'Not a clue, but it works, doesn't it?' (2007, p. 255; 2017, p. 350).

Heterophenomenological agnosticism obviously makes sense for the new Chicago reports; and it should just as obviously make sense for our reports about conscious experience. All the experiences we take for granted are just as strange to us as this one; we think we have much more access to our own experiences than we can ever convey to other people through verbal report, but we do not. We are therefore deluded if we think that autophenomenology (the study of one's own phenomenology) is a 'more intimate, more authentic, more direct way' of studying consciousness than heterophenomenology (the study of another person's phenomenology) (2017, p. 351). And we should not take the immense variety of introspective reports between individuals (including on questions as fundamental as whether or not thought has a distinctive phenomenology) as evidence that our conscious experiences really are vastly divergent, says philosopher

Eric Schwitzgebel (2008): that variety is just more evidence for the profound unreliability of introspection.

Taking a defamiliarising stance on your own experience is something we have tried to encourage with the Practices throughout this book. Not taking your own 'what-it's-like' at face value is an important habit to get into if you want to take your exploration of consciousness further.

Another good habit is to drop your defensiveness towards other disciplines. Whether your home field feels like psychology or neuroscience or philosophy, particle physics or literary studies, it is clear that the mystery of consciousness is not going to be solved any time soon by any single existing disciplinary paradigm. It's very easy to be protectionist, perhaps especially if we have grown up in the humanities and sometimes feel a little resentful that science gets all the status and all the money. But protectionism usually leads to misreading, caricature, and many missed opportunities for exciting research. And there's nothing that really needs protecting – certainly not 'the citadel of the first-person' (Dennett, 2007, p. 264).

Heterophenomenology may be a good way of proceeding – a good default for the moment, while we work out where and how to safely use phenomenological reports (van de Laar, 2008). Indeed, the essence of heterophenomenology is simply not committing ourselves in advance. Part of this is waiting until we know more: 'It sure seems as if there is a Cartesian Theater. But there isn't. Heterophenomenology is designed to honor these two facts in as neutral a way as possible until we can explain them in detail' (Dennett, 2007, p. 269). For Dennett, the most promising 'for now' attitude is to rephrase the mystery of consciousness in the way Newton rephrased the mystery of gravity: to stop asking what it is (a fluid, a substance, a force?) and start asking how it behaves. For Dennett, phenomenologists are in practice committed to this 'bland form of behaviourism' without realising it, because like everyone else, they have nothing else to say about consciousness beyond what it does.

But in a flash, as of lightning, all our explanations, all our classifications and derivations, our aetiologies, suddenly appeared to me like a thin net. That great passive monster, reality, was no longer dead, easy to handle. It was full of a mysterious vigour, new forms, new possibilities. The net was nothing, reality burst through it. [. . .] That simple phrase, I do not know, was my own pillar of fire. For me, too, it brought a new humility akin to fierceness. For me too a profound mystery. [. . .] There had always been a conflict in me between mystery and meaning. I had pursued the latter, worshipped the latter as a doctor. As a socialist and rationalist. But then I saw that the attempt to scientize reality, to name it and categorize it and vivisect it out of existence, was like trying to remove the air from the atmosphere. In the creating of the vacuum it was the experimenter who died, because he was inside the vacuum.

(John Fowles, *The Magus*, 1965/2010, p. 309)

'Are we not familiar enough with our own experiences?'

(Gray, 2004, p. 123)

'the widespread conviction that you have to defend the citadel of the first-person is simply a mistake'

(Dennett, 2007, p. 264)

'Self-deception may feel like insight'

(Metzinger, 2009, p. 220)

Maybe a final part of the complicated jigsaw of different ways of studying consciousness is that we are scared of what we might find if we abandon the methods that are familiar to us. We may well have 'quite reasonable anxieties about whether we might hate what we eventually learned about our own brains and mind, and these anxieties promote wishful thinking *on all sides*' (Dennett, 2007, p. 269). It takes courage to set aside what you think you know, is this is nowhere more true than of the experiences that feel so intimately yours. In the final chapter, we will hear more from people who have trained with great commitment in noncommittal self-observation. What do they have to say about what they find?

READING

Dennett, D. C. (2007). Heterophenomenology reconsidered. *Phenomenology and the Cognitive Sciences*, 6(1-2), 247-270.

A response to fifteen peer commentaries in the same special issue on different versions of phenomenology (and other ways of studying consciousness).

Garrison, K. A., Santoyo, J. F., Davis, J. H., Thornhill, T. A., Kerr, C. E., and Brewer, J. A. (2013). Effortless awareness: Using real-time neurofeedback to investigate correlates of posterior cingulate cortex activity in meditators' self-report. *Frontiers in Human Neuroscience*, 7, article 440.

Uses grounded theory to analyse self-reports and derive distinctions between different meditative states during real-time neurofeedback sessions.

Petitmengin, C., and Lachaux, J.-P. (2013). Microcognitive science: Bridging experiential and neuronal microdynamics. *Frontiers in Human Neuroscience*, 7, article 617.


Suggests new methods to tackle the problem of different levels of description between neurophenomenology and neuroimaging.

Price, D. D., and Barrell, J. J. (2012). Developing a science of human meanings and consciousness. In D.D. Price and J.J. Barrell, *Inner experience and neuroscience: Merging both perspectives* (pp. 1–30). Cambridge, MA: MIT Press.

An overview of possible methods for a science of meaning and experience, focusing on what to do with the first person.

Thompson, E. and Zahavi, D. (2007). Phenomenology, In P.D. Zelazo, M. Moskovich, and E. Thompson (Eds), *The Cambridge handbook of consciousness* (pp. 67–87). Cambridge: Cambridge University Press.

An introduction to the past and future of phenomenology's methodological and theoretical approaches to consciousness.



Waking up

CHAPTER

EIGHTEEN

Once upon a time, about two and a half thousand years ago, a prince was born in the north of India. His name was Siddhartha Gautama, and he led a happy and indulgent childhood, protected from the harsh realities of life. One day he walked out of his comfortable palace into the streets and saw a sick man, an old man, a beggar, and a corpse. Shocked by all the suffering he saw, and the contrast with his own life, he vowed to search for the meaning of existence. When he was 29 he left behind his wealth, his wife, and his young son, and set off to become a wandering ascetic, depriving himself of every comfort and outdoing all the other ascetics of his time with harsh self-imposed discipline. After six years, when almost starving to death, he accepted some milk gruel, gradually regained his health, and concluded that neither indulgence nor deprivation was the way to truth: a middle way was needed. He sat down under a pipal tree and vowed not to get up again until he understood.

On the seventh day, with the morning star shining in the sky, he became enlightened. This famous story, based on at least some historical fact, is the tale of an ordinary person waking up – and this is how the Buddha described what had happened. He had woken up.

He realised that what he had seen was there for all but could not be spoken about. So he could not see how to teach. Yet people flocked to him and so he spent the next forty-five years travelling widely and teaching the 'Four Noble Truths' and other teachings that became known as the dharma. He urged people not to be satisfied with hearsay or tradition but to look within to see the truth, and it is said that his last words were 'Work out your own salvation with diligence.'

'Zen can no more be explained than a joke. You see it or you don't'

(Humphreys, 1949/1951, p. 3)

The four noble, or ennobling, truths are 1) *Suffering*. Life is inherently unsatisfactory because everything is impermanent. 2) *The origins of suffering*. We suffer because we cling to things we like and reject those we don't, becoming trapped in a cycle of being and becoming called *samsara*. 3) *The cessation of suffering*. Recognising impermanence, and letting go of desire and the desiring self, ends suffering. Sorrow and grief, joy and happiness can come and go without attachment, leading to *nirvana*. 4) *The way*. The Buddha recommended an eightfold path of right understanding, thought, speech, action, livelihood, effort, mindfulness, and concentration. Heart and mind, or wisdom and compassion, are seen as inseparable.

This path has been described as

a way of life to be followed, practised and developed by each individual. It is self-discipline in body, word and mind, self-development and self-purification. It has nothing to do with belief, prayer, worship or ceremony. In that sense, it has nothing which may popularly be called 'religious'.

(Rahula, 1959, pp. 49–50)

Nevertheless, many of the Buddha's discourses were passed on by oral tradition and then written down several hundred years later as the *sutras*, and despite his warnings about relying on hearsay and tradition, Buddhism became a great religion, spreading to southern India, Ceylon, and Burma as Theravada Buddhism, and elsewhere as Mahayana Buddhism. It spread to Tibet, where it took a unique form built on existing folklore, including reliance on the concept of reincarnation, which was already popular there. It spread along the silk route from India to China, where it became Chan Buddhism, to Japan where Chan became Zen, and eventually also to the West (Humphreys, 1951; Batchelor, 1994).

In this final chapter, we will explore some of the questions that remain after our long journey through the wide and often baffling territories of research and thought that relate to consciousness. This chapter is more personal than the other seventeen, but we hope that some of your questions might be similar to ours, or that ours will interest you nonetheless. The questions we will try to address are as follows.

Are we stuck with the problems and illusions we have discovered, or can we learn to see through them? If we can, this amounts to wondering: can consciousness itself change?



PROFILE 18.1

Sam Harris (b. 1967)



Trained as both philosopher and neuroscientist, Sam Harris has a PhD in cognitive neuroscience from UCLA, has written books on faith, religion, morality, and free will, and runs the *Waking Up* podcast. He is fiercely critical of organised religion

and along with Dennett and Dawkins is considered one of the 'Four Horsemen' of the new atheism, although, unlike them, he is a long-term meditator, believing that some Buddhist and Hindu traditions offer valuable empirical insights into consciousness. Experiences with psychedelic drugs, including LSD, psilocybin, and MDMA, led him to leave Stanford in his second year to seek spiritual insight without drugs. Travelling to India, he pursued strenuous methods of meditation, including a year on silent retreat, concluding that the key aim is to look into the sense of being a separate self until it dissolves. He thinks free will is an illusion, morality can be studied scientifically, everything we do is for the purpose of altering consciousness, and Brazilian Jiu Jitsu (though having little to do with consciousness) is surprisingly relevant to the illusoriness of the ego.

We began this chapter with the Buddha as an example of a person who claimed that such transformation is possible. Since then, many secular people have claimed the same thing. Many people describe the change to their consciousness as something like waking up. Does their consciousness really change, and if so, can we find out what kind of change is possible, and what its consequences may be?

We began this book with the warning that learning about consciousness will change your life. Here we restate that warning as a question, and try to start answering it.

We will talk quite a lot about Buddhism, because it is one of the contexts where spiritual and scientific learning have come closest to one another. But neither of us is a Buddhist, and if you are not, neither of us wants you to become one. We do not equate spirituality with religion, but we accept that they have long been intertwined. We are interested in the mind, and in what personal practice can and cannot change.

Once Zhuang Zhou dreamt he was a butterfly, a butterfly flitting gaily. He knew nothing of Zhou. Suddenly, he awoke, and all at once he was Zhou. But he didn't know whether Zhou had dreamt he was a butterfly or a butterfly was dreaming he was Zhou. Surely there is a difference between Zhou and a butterfly – this is what we call the transformation of things!

(‘The butterfly dream’, third century BC, translated by Robert Eno (2010/2016), p. 23, *Zhuangzi: The Inner Chapters*; see also Thompson, 2014, pp. 198–202, for reflections on this parable)

BUDDHISM IN SCIENCE

Science and religion are often opposed, not least because most religions rely on unchanging sacred books and teachings, while science constantly updates itself, seeking to understand the world by interrogating it with experiments. Yet Buddhism has found a place within psychology in a way no other religious teachings have. There have been many books and conferences on East–West psychology from the 1980s onwards, and the vast majority of contributions have dealt with Buddhism rather than other traditions (Claxton, 1986b; Crook and Fontana, 1990; Lama et al., 1991; Pickering, 1997; Watson et al., 1999; Segal, 2003; Hanson and Mendius, 2009). In 1987 the current Dalai Lama, head of the Gelug sect of Tibetan Buddhism, began a series of dialogues with Western scientists, and in 2005, despite some protests that this was an inappropriate mixing of science and religion, was invited to speak at the annual Society for Neuroscience conference. He has since continued trying to build bridges between Buddhism and science. These efforts have resulted in projects like the Atlas of Emotions compiled by psychologist Paul Ekman and social science and public health researcher Eve Ekman in collaboration with the Dalai Lama, intended as a tool to help people choose, if not what emotions to feel, at least how to respond to those we do feel.

There are many possible reasons for this dialogue. Unlike Christianity, Judaism, and Islam, Buddhism has no god, no supreme creator, and no notion of an indestructible human soul. In his book *Waking Up*, American neuroscientist Sam Harris

compares the Buddha with Jesus. While the historical Buddha, Siddhartha Gautama, was 'merely a man who woke up from the dream of being a separate self', Jesus was supposed to be the son of the creator of the universe. This, says Harris, 'renders Christianity, no matter how fully divested of metaphysical baggage, all but irrelevant to a scientific discussion about the human condition (2014, p. 30).

Hinduism shares with Buddhism the idea that we live our ordinary lives in the world of *maya* or illusion: an unenlightened dream of duality in which self and the universe seem to be distinct. But most of its many traditions also include personal and celestial deities, and the idea that each of us is or has an eternal self or soul called the *atman*. Nonetheless, the highest principle of Hindu philosophy, *Brahman*, is not a personified deity, but an impersonal spiritual force, the ultimate reality of the universe, and there are nondualist traditions in Hinduism, especially in Advaita, in which Brahman and atman are ultimately found to be identical – although both may still be considered distinct from a material, bodily reality. Buddhism is more fully atheist and teaches no-self and nonduality more consistently than any other tradition.

Buddhists are also urged not to worship anyone or believe any doctrines, but to inquire into their own minds, and have faith that they too can wake up. In *Buddhism without Beliefs*, the British scholar and Zen Buddhist Stephen Batchelor (1997) explains that the noble truths are not propositions to be believed in, but truths to be acted upon. Practising Buddhism entails an inquiry into oneself that supposedly reveals the emptiness and impermanence of all phenomena, the illusory nature of self, and the origins and ending of suffering. Harris agrees. The teachings of Buddhism are, he says, 'empirical instructions: If you do X, you will experience Y' (2014, p. 30). This is reminiscent of Max Velmans's motto for all of science: 'If you carry out these procedures you will observe or experience these results' (1999, p. 300).

This structural affinity with science runs deep within Buddhist teachings. A central teaching in all branches of Buddhism is the doctrine of conditioned arising, or co-dependent origination. The Buddha taught that all things are relative and interdependent, arising out of what came before and in turn giving rise to something else in a vast web of causes and effects. This can be seen as a very early statement of a scientific principle of cause and effect – and of the conviction that there is no magic involved, no skyhooks. Not accepting this is one source of illusion, or ignorance. This principle is applied specifically to consciousness as well as to everything else, and the Buddha denied the possibility of there being consciousness without the matter, sensations, perceptions, and actions that condition it (Rahula, 1959). This conception of an interconnected, causal universe is compatible with basic physics and modern science in a way that a universe created and sustained by any god is not.

A good example of this is the practice of the Buddhist jhanas, the series of eight increasingly absorbed states reached through deep concentration (Chapter 13). Although they are rarely practised these days, the American meditation teacher Leigh Brasington (2015) gives precise instructions for inducing them. After a series of preliminary concentration practices, the first involves concentrating on positive emotions and feelings throughout the body. This induces what feels like a flood of energy, manifesting in heat, shaking, or trembling that sounds rather like the esoteric notion of a kind of 'primal energy' called kundalini. This is then

modulated by further steps leading to the next states in the series. Far from invoking occult ideas, Brasington speculates that these states may depend upon a cascade of several neurotransmitters, and he takes part in research designed to link the unfolding sequence of altered states with the underlying physiology (Hagerty et al., 2013). He teaches these skills very much in the tradition of ‘If you do X, you will experience Y’, holding out the hope that future research will uncover the reasons why doing X leads reliably to experiencing Y.

American philosopher William Mikulas agrees that an important reason for the dialogue between science and Buddhism is that it focuses on methods not doctrines. He describes essential Buddhism as having ‘no creeds or dogmas, no rituals or worship, no saviour, and nothing to take on faith; rather it is a set of practices and free inquiry by which one sees for oneself the truth and usefulness of the teachings’ (2007, p. 6). Many other scholars make similar points and, like Mikulas, refer to ‘essential Buddhism’ as though this can easily be extracted from all the later accretions and different sects, but we should note that in many parts of the world Buddhism is as much involved in rituals and belief systems as any other religion. Even so, his point was that ‘The Buddha made no claims about himself other than that he woke up. [. . .] The possibility and nature of such awakening is a major challenge to North American academic psychology’ (p. 34). This is a challenge that has since been enthusiastically taken up (Hanson, 2009; Michaelson, 2013; Taylor, 2017).

A pertinent question is: what happens if your earnest inquiry into yourself provides answers other than those about impermanence, illusion, and suffering? What if your waking-up is different from the Buddha’s? Many writers have described how ordinary people, as well as devoted practitioners of meditation, have just ‘woken up’, and their accounts include again and again the familiar notions of freedom from illusion and the ending of duality that lightens suffering (Kapleau, 1980; Crook and Fontana, 1990; Sheng-Yen et al., 2002; Harris, 2014). Does this mean that waking up is always the same? Not necessarily. Those practising within any tradition will inevitably be influenced by their teachers, and the effects of meditation may be heavily dependent on their expectations. Even if people spontaneously wake up with no knowledge of spiritual or mystical experiences, they may simply be falling for a common illusion – and we have met plenty of examples of such common illusions. How can we be sure that claims to have dropped the illusions are to be believed? This is a question that the scientific study of spiritual experience must address, and indeed is beginning to address.

There are other reasons to be sceptical about the fit between Buddhism and science. Some of the core Buddhist teachings, such as the *Abhidharma*, may appear to be more like psychology than doctrine, including complex categorisations and long lists of mental phenomena with their origins and interconnections. Yet, unlike a scientific psychology, these schemes are fixed and unchanging, more akin to doctrines to be learned and believed than hypotheses about the mind to be tested by any kind of experimentation. In this sense, they become more like religious dogma than the Buddha’s urging not to depend on scriptures and doctrines but to work towards one’s own awakening.

There may seem to be another difference from conventional science in that the *Abhidharma’s* categories of mind are derived not from third-person experiments but from first-person phenomenological inquiry, but in the previous chapter we

concluded that the gap between the two may not be as wide as it seems. Indeed, Varela suggested that Buddhist mindfulness meditation could be used in neurophenomenology, and that 'the Buddhist doctrines of no-self and of nondualism that grew out of this method have a significant contribution to make in a dialogue with cognitive science' (Varela, Thompson, and Rosch, 1991, p. 21).



PRACTICE 18.1

WHAT IS THIS?

Read the story about Hui Neng and the monk in Concept 18.1. Think about the question he asked: '**What is this thing and how did it get here?**' Think about it as applied to the monk, standing there at the monastery after days of walking in the mountains. Think about it as applies to yourself, sitting here, walking there, realising you haven't thought about the question for half an hour, and now standing here. Think about it whatever you are doing. 'What is this thing and how did it get here?' Go on asking the question all the time. The words do not matter. As you carry on practising, they will probably fall away until you begin the question and 'Wh. . .?'

These may be some of the reasons why many psychologists have turned to Buddhism and found both methods and theories relevant to the psychology of consciousness. A large proportion of these focus on the Zen tradition within Buddhism. Why? Because, according to American neurologist and author James Austin, Zen is 'the approach most systematic yet most elusive, the clearest yet most paradoxical, the subtlest yet most dramatic' (1998, p. 7), and is 'untainted by belief in the supernatural or the superstitious' (Kapleau, 1980, p. 64). It is also less preoccupied with outward forms than Tibetan Buddhism, which uses elaborate altars and images, and complex visualisations of deities, each with different



FIGURE 18.1 • A young monk arrives at Hui-Neng's famous mountain monastery. What is it?

movements, clothes, adornments, and colours. These elaborate techniques can be very powerful for inducing ASCs and training concentration and attention, but they do not necessarily appeal to philosophers and scientists seeking to understand the mind.

Then there is the vexed question of reincarnation. This is prominent in Tibetan Buddhism, which was grafted onto existing folk beliefs in reincarnation, but less so in Zen, which developed in China and Japan. The popular conception of a personal reincarnation in which some lasting essence passes through many lives seems to make no sense to the Western scientific mind. Indeed, it makes little sense within the context of the Buddha's teaching of the impermanence and emptiness of self, for what is there to be reincarnated?

Zen has a tradition of avoiding most of this and going straight to the point. 'Zen is the apotheosis of Buddhism', says Christmas Humphreys, the founder of the Buddhist Society in Britain.

This direct assault upon the citadel of Truth, without reliance upon concepts (of God or soul or salvation), or the use of scripture, ritual or vow, is unique. [. . .] In Zen the familiar props of religion are cast away. An image may be used for devotional purposes, but if the room is cold it may be flung into the fire.

(1951, p. 179–180)

'The object [of attention] determines whether a meditation practice is religious, therapeutic, or something else'

(Mikulas, 2007, p. 24)

The real task in hand is that 'the mind may be freed'.

This comment raises the question of how the objectives of science and Buddhism compare. Freeing the mind could be understood as an ambition compatible with science and philosophy: a free mind can find out truth for the sake of truth. But Buddhism is usually thought of as trying to find out the truth in order to transform oneself, to become free from suffering, and even to save all sentient beings from suffering. In this sense, Buddhism may be closer to psychotherapy than to science.

TRANSFORMATION AND THERAPY

In a story from Tibetan Buddhism, a poor, low-caste woodcutter called Shalipa lived near the charnel ground where corpses were thrown to rot. Shalipa was so terrified of the corpses and the wolves howling at night that he couldn't eat or sleep. One evening, a wandering yogin came by asking for food and Shalipa begged him for a spell to stop the howling. The yogin laughed, 'What good will it do you to destroy the howling of the wolves when you don't know what hearing or any other sense is. If you will follow my instructions, I will teach you to destroy all fear.' So Shalipa moved inside the charnel ground and began to meditate upon all sound as the same as the howling of wolves. Gradually he came to understand the nature of sound and of all reality. After nine years he lost all fear, attained great realisation, and became a teacher himself, wearing a wolf skin around his shoulders.

Shalipa is just like us, says American psychologist Eleanor Rosch (1997), even though he lived so long ago and so far away. There he is, shivering in his hut with all his social, psychological, medical, and spiritual problems. This is the common state and is why our modern psychology is based on such a dualistic and alienated view of the human condition. In Buddhism this deluded state is called *samsara*,

the idea that *you* are trapped in the wheel of birth and death: enlightenment is freedom from *samsara*. So the yogin does not advise Shalipa to sue the owners of the charnel ground, to delve into the meaning of wolf howls in his personal history, or to endure his fate to obtain religious salvation. He teaches him to use his own experience as a means of radical transformation. The new Shalipa has no fear because he is free of illusion.

In the meeting between Buddhism and psychotherapy, one live question concerns whether, fundamentally, the two endeavours are the same (Claxton, 1986b; Pickering, 1997; Watson, Batchelor, and Claxton, 1999; Mikulas, 2007). The British philosopher Alan Watts (1961) brought Eastern teachings to the West in the 1940s and wrote extensively on Zen. He said that looking into Buddhism, Taoism, Vedanta, and Yoga, we do not find either philosophy or religion as these are understood in the West; we find something more nearly resembling psychotherapy. Even so, he pointed out many differences, not least in the lengths of their traditions and their different responses to the problem of suffering.

Although both aim to transform the individual, their methods are strikingly different, and so is the kind of transformation they seek. While psychotherapy aims to create a coherent sense of self, Buddhist psychology aims to transcend the self. Types of therapy differ widely, but broadly speaking they all aim to improve people's lives and to make them healthier, happier, and less anxious. So a successful outcome for most therapy is a person who is happy, relaxed, well-adjusted to their society, and able to function well in their relationships and occupation. A successful outcome for a Buddhist might be the same, but it might equally be a hermit who shuns all society and lives in a cave, a teacher who rejects all conventional teachings, or a wild and crazy wise one whose equanimity and compassion shines through their mad behaviour.

Claxton suggests that therapy is a special and limited case of the more general spiritual search. While therapists and clients may agree to leave certain useful defences in place, on the spiritual path nothing is left unquestioned. 'The quest is for Truth not Happiness, and if happiness or security or social acceptability must be sacrificed in the pursuit of this ruthless enquiry, then so be it' (1986b, p. 316). For John Crook (1980), Zen training is more like 'total therapy' in which the cage of identity is broken. This may be so to the extent that self is a cage, but Mikulas warns of a common misunderstanding among spiritual practitioners, 'that one must undo or kill the personal level self in order to awaken; but this is not necessary or desirable' (2007, p. 34). Awakening is not about eliminating or devaluing the self; it is about letting go of identifying with the self, or losing the sense that self is separate from the rest of the universe. And this is what leads to freedom and peace of mind.

For some, the spiritual enterprise takes off where therapy ends, implying that psychotherapy must come before the greater task of seeing through the self. This suggests a developmental or 'full-spectrum' model of consciousness leading not only from infancy to adulthood, but from immaturity to full enlightenment. There have been several attempts to develop such models, including the complex multilevel schemes proposed by the Buddhist writer Ken Wilber (Wilber, Engler, and Brown, 1986; Wilber, 2001) and the ideas of American psychotherapist and Buddhist Jack Engler.

'you have to be somebody before you can be nobody'

(Engler, 1986, p. 49)

Engler famously proclaimed, 'you have to be somebody before you can be nobody' (1986, p. 49). He studied the effects of Buddhist practice on students with a wide range of different starting points, and found that those who were attracted to Buddhism because of failures in self-development or as a way of avoiding facing themselves ran the risk of further fragmenting their already fragile sense of self. This suggested that an early emphasis on self-transcendence might be therapeutically harmful and therefore not help the transcendental aim, either. He concluded that a sense of both self and no-self are necessary, and in that order. He later changed his view somewhat, stressing that our motivations and our conflicts are so complex that a neat developmental model doesn't quite work. But he continued to argue for the importance of "being somebody" – that is, facing crucial developmental or life stages head on instead of attempting to avoid them in the name of spirituality or enlightenment' (2003, p. 36). This self-related task is important even though 'the experience of being or having a self is a case of mistaken identity, a misrepresentation born of anxiety and conflict about who I am' (p. 36).

Interweaving therapy with Buddhism is now fairly common, in various strands of transpersonal psychology, and in forms of psychotherapy that include Buddhist methods of practice. Some people argue that although psychological and spiritual work address different levels of human existence, spiritual work can have therapeutic value, and therapeutic methods can help in the integration of spiritual insights into ordinary life (Watson, Batchelor, and Claxton, 1999). Examples include Kabat-Zinn's MBSR (Chapter 7), which emphasises paying attention and developing a non-judging awareness in order to break through the 'unconscious consensus trance that we think of as being awake' (Kabat-Zinn, 1999, p. 231). Crook integrated therapeutic techniques into his 'Western Zen' retreats (Crook and Fontana, 1990), and breathing techniques, mindfulness, and meditation have been frequently used in schools, prisons, sports, parenting, and many other contexts (Watson, Batchelor, and Claxton, 1999). Mindfulness-based cognitive therapy is an eight-session group intervention programme based on MBSR, designed by Mark Williams, Zindel Segal, and John Teasdale, and meta-analyses have shown that it can reduce risk of relapse to depression for those who have experienced three or more episodes (Piet and Hougaard, 2011). MBCT also seems to be about

as effective as standard cognitive behavioural therapy at reducing anxiety, depression, and stress, with mindfulness levels increasing during the programme and strongly correlating with clinical outcomes (Khoury et al., 2013). (But see the website for discussion of the difficulties of doing conclusive research in this area.)

Those who persevere with spiritual practice claim many therapeutic effects, in particular that they become more loving, compassionate, and equanimous. It may seem odd that letting go of desire, giving up your self, and treating everything as impermanent can possibly have such effects. Surely, goes the worry, if you stop controlling yourself terrible disasters will ensue

(Levine, 1979; Rosch, 1997). This is the same fear that attends the idea of giving up free will, and indeed giving up the sense of being, or having, a separate self does do away with the feeling of being in control or of having free will. Yet, as



FIGURE 18.2 • The headless view. To others you are a person in the world. To yourself you are a space in which the world happens.

we learned in Chapter 9, Claxton did not end up running over old ladies for fun and Harris felt his ethics and his state of mind had improved, not deteriorated. Nonetheless, it would be foolish to embark on a search for spiritual transformation expecting it to make you happier. It may, or it may not. But, as many traditions point out, chasing after happiness may itself get in the way of finding it.

SPONTANEOUS AWAKENING

Awakening is often described as though it were the endpoint of a long journey on a spiritual path, but some people claim that they just *woke up*, and that their awakening was the beginning, rather than the culmination, of their spiritual life.

The best day of Douglas Harding's life, his rebirth-day, as he called it, was when he found he had no head. At the age of 33, during the Second World War, he had long been pondering the question 'What am I?' One day, while walking in the Himalayas, he suddenly stopped thinking and forgot everything. Past and future dropped away, and he just looked. 'To look was enough. And what I found was khaki trouserlegs terminating downwards in a pair of brown shoes, khaki sleeves terminating sideways in a pair of pink hands, and a khaki shirtfront terminating upwards in – absolutely nothing whatever!' (Harding, 1961, p. 2).

We can all do what he did next. We can look where the head should be and find a whole world. Far from being nothing, the space where the head should be is filled with everything we can see, including the fuzzy end of our nose and the whole world around. For Harding, this great world of mountains and trees was completely without 'me', and it felt like suddenly waking up from the sleep of ordinary life. It was a revelation of the perfectly obvious. He felt only peace, quiet joy, and the sensation of having dropped an intolerable burden.

Harding stresses that headlessness is obvious if you look clearly. There are not two parallel worlds, an inner and an outer world, because if you really look you just see the one world which is always before you. This way of looking explodes the fiction of inside and out, and of the mythical centre; it explodes 'this terminal spot where "I" or "my consciousness" is supposed to be located' (1961, p. 13). He might equally have said that it blows up the Cartesian theatre.

Harding soon discovered that others did not share his revelation. When he tried to explain it, people either thought he was mad or said 'So what?', but eventually he stumbled upon Zen. There he found others who had seen as he did, such as Hui Neng, who told a fellow monk to see. 'See what at this very moment your own



ACTIVITY 18.1

The headless way

Here are two little tricks to do all together in class, or on your own. Some people can be flipped into an entirely new way of experiencing, but others just say 'So what'. So the tricks may, or may not, work for you. Take them slowly and pay attention to your own immediate experience. Don't rush.

Pointing. Point at the window, and look carefully at what you see there. Note both your finger and the place it points at. Point at the floor, and look carefully at where your finger is pointing. Point at your foot, and look carefully at what you see. Point at your tummy, and look carefully at what is there. Point at yourself, and look carefully at what you see there.

What did you find there?

Head to head. Find a friend to work with. Place your hands on each other's shoulders, and look steadily at your friend's face and head. Now ask yourself – how many heads are there? Don't think about what you know, or what must be true; pay attention to your own direct experience now. How many heads can you see? What, in this present experience, is on the top of your shoulders?

'It's not a doing but an undoing, a giving up, an abandonment of the false belief that there's anyone here to abandon. What else is there to do?'

(Harding, 1961, p. 73)

'It is possible to stand free of the juggernaut of self, if only for moments at a time'

(Harris, 2014, p. 11)

face looks like – the face you had before you (and indeed your parents) were born.' This became one of the most famous of Zen koans and was exactly to the point.

Even so, many continue to reject Harding's simple insight. Hofstadter calls it 'a charmingly childish and solipsistic view of the human condition. It is something that, at an intellectual level, offends and appalls us' (Hofstadter and Dennett, 1981, p. 30). He seems unable to imagine that, as Harris puts it, 'It is possible to stand free of the juggernaut of self, if only for moments at a time' (Harris, 2014, p. 11); unable to imagine the state that Harding describes, and that so many others have too, of being alert and alive yet utterly without the sense of an observing self.

John Wren-Lewis was a physics professor with decidedly anti-mystical views when in 1983, at the age of 60, he was poisoned while travelling on a bus in Thailand. A would-be thief gave him a toffee laced with what was probably a mixture of morphine and cocaine, and the next thing he knew was waking up in a dilapidated and dirty hospital.

At first he noticed nothing special, but gradually it dawned on him that it was as if he had emerged freshly made, complete with the memories that made up his personal self, from a radiant vast blackness beyond space or time. There was no sense at all of personal continuity. Moreover, the 'dazzling darkness' was still there. It seemed to be behind his head, continually recreating his entire consciousness afresh, instant by instant, now! and now! and now! He even put his hand up to feel the back of his head only to find it perfectly normal. He felt only gratitude towards everything around him, all of which seemed perfectly right and as it should be.

Both doctors and patient thought that the effects would soon wear off, but they did not, and years later Wren-Lewis described how his whole consciousness had changed for good.

I feel as if the back of my head has been sawn off so that it is no longer the 60-year-old John who looks out at the world, but the shining dark infinite void that in some extraordinary way is also "I".

(1988, p. 116)

Many aspects of his life changed. The practicalities of ordinary life became easier, not harder as you might imagine, because he was not constantly



CONCEPT 18.1

KOANS

Working with a koan or hua-tou is a method used to induce deep Zen questioning, originally developed in China from the sixth century onwards. Among famous collections are the 100 koans of the *Blue Cliff Record*, compiled in 1125 and later expanded by many commentators, including the poet and painter Hakuin (1685–1768), and the *Gateless Gate* collection of 48 koans devised by Wumen in 1228. Koans are mainly used in Rinzai Zen, one of the two main sects, and trainee monks may be expected to 'pass' a series of graded koans, but really koans do not have 'right answers'. The only right answer is to show that one has 'seen the nature' or 'transcended duality' (Watts, 1957; Kapleau, 1980).

Many koans are questions directed at the nature of self, such as 'What was your original face before your mother and father were born?', 'What is your own mind?', or 'Who is dragging this old corpse around?' It is easy to spend hours, days, or even years on any of these. If you have been doing the practices in each chapter, you will know just what this means. Indeed, Sue has used some of these as the basis for prolonged meditation, including 'Am I conscious now?' and 'Who is asking the question?' (Blackmore, 2011). Other koans may seem completely incomprehensible, such as 'The East Mountain strides over the water' or 'When the many are reduced to one, to what is the one reduced?', yet they may have deep effects on the serious questioner.

thinking about the future. Pain became more of an interesting warning sensation than a form of suffering. His sleep changed from a previously rich dream life to a state of 'conscious sleep' in which he was still aware of lying in bed, and the fifty-nine years of his former life seemed like a kind of waking dream. He was no longer living with an illusion of separate selfhood; rather, everything had become 'just the universe John Wren-Lewis' (2004).

His original experience might be classed as an NDE (Chapter 15), yet Wren-Lewis came to precisely the opposite conclusion from that of most NDE researchers. Rather than leaping to ideas about consciousness existing apart from the brain or notions of 'endless consciousness' or overthrowing reductionism (Haesler and Beauregard, 2013; van Lommel, 2013; Parnia et al., 2014), he concluded that his personal consciousness was 'snuffed out' and then recreated from the radiant dark.

This is more reminiscent of Dan Dennett's (1991) suggestion that self and consciousness can always be snuffed out like a candle flame and rekindled later, or of James's passing Thoughts, or Galen Strawson's sense that consciousness is continually restarting. Wren-Lewis was acutely aware of that snuffing out and re-creation going on all the time. This is the same idea as the Buddhist notion of the wheel of death and rebirth – the constant re-creation of new selves moment-to-moment from the prevailing conditions that we mistake for a continuing self. Our problem is that we don't tend to see it that way; instead, we want our self to be permanent.

As for the spiritual path, Wren-Lewis claimed that the very idea is necessarily self-defeating, because the process of seeking implies a preoccupation with time, and so makes a goal out of what is already here and now. In this, he is expressing the paradox of the path to no-path found so often in Zen. He is particularly scathing about those philosophies which are based on schemes of spiritual growth or conscious evolution. Awakening is not the culmination of a journey but the realisation that you never left home and never could.

These examples show, unequivocally, that awakening does not have to be the culmination of a long process of training. Harding woke up through lone questioning and happenstance, and Wren-Lewis through a poisoned brain. Research suggests that

Hua-tous are the head, or final words or questions, of a Zen story. A famous example forms the basis of Korean Zen (S. Batchelor, 1990; M. Batchelor, 2001). It comes from the turn of the eighth century, when it was common for teachers to point to a house, or the sky, or a leaf, and demand 'What is that?' As the story goes, a young monk walked for many days to find the Zen patriarch Hui Neng at his mountain monastery. The ragged monk who met him at the gate chatted politely about his journey and then demanded, 'What is this thing and how did it get here?' Not realising this was indeed Hui Neng himself, the monk was speechless and decided to stay and devote himself to this question.

After eight years of practice he finally went to Hui Neng again and said, 'I have experienced some awakening'. 'What is it?' asked Hui Neng. The monk replied, 'To say it is like something is not to the point. But still it can be cultivated'.

Using this koan means that when walking, standing, sitting, or lying down, you repeatedly ask the question 'What is this?', meaning 'What is walking?', 'What tastes the tea?', or 'What is it before you even taste the tea?' With practice you do not need to repeat the words; it is the doubt or perplexity that matters. So the question hangs there, always being asked. Your whole body and mind become the question. You don't know.

Here is a different view. Hofstadter (2007) has Strange Loop #641 ('a somewhat churlish proxy for the author', p. 277) describe a koan:

A master was asked the question, 'What is the Way?' by a curious monk. 'It is right before your eyes', said the master. 'Why do I not see it for myself?' 'Because you are thinking of yourself.' 'What about you – do you see it?' 'So long as you see double, saying "I don't" and "you do" and so on, your eyes are clouded.' 'When there is neither "I" nor "You", can one see it?' The master replied, 'When there is neither "I" nor "You", who is the one that wants to see it?'

Strange Loop #641 calls it 'Just a bunch of non sequiturs posing as something that should be taken with the utmost gravity' (p. 300).

‘But the deepest goal of spirituality is freedom from the illusion of the self – and to seek such freedom [. . .] is to reinforce the chains of one’s apparent bondage in each moment’

(Harris, 2014, p. 123)

‘the brief awakenings of kensho and satori are “nothing special”’

(Austin, 2009, p. 111)

‘does ordinary insight differ from the brief state of “enlightenment” called kensho in Zen?’

(Austin, 2009, p. 125)

permanent psychological transformation can occur suddenly in response to turmoil or trauma, perhaps because the intensity of suffering means that psychological attachments have to be dissolved, and acceptance ensues (Taylor, 2012). It may also be the case that trying hard to get good at introspection only reinforces the sense of self that drops away at other times (Goldberg et al., 2006; discussion in Block, 2007, and commentaries). But this does not mean that training and practice are useless. Perhaps deep questioning can prepare the person in some way. On the other hand, perhaps poisoning can change a brain in ways comparable to long years of meditation. Or perhaps there is just an element of luck about it all. As one contemporary aphorism has it, enlightenment is an accident, and meditation helps you become more accident-prone.

ENLIGHTENMENT

What, then, is enlightenment? Although Shakyamuni Buddha’s story is often taken as the apotheosis of enlightenment, there were probably many people before him who went through this transformation, and many are said to have done so since. Among them are old and young, men and women, monks and laypeople. They include modern Westerners: businessmen, artists, homemakers, and psychologists (Kapleau, 1980; Sheng-Yen et al., 2002; Harris, 2014). We give it a name, and its name – in Buddhist terminology, *bodhi* – means something like awakening. Yet it is hard to say what has happened to them. Enlightenment is apparently a profound transformation of what it’s like to be. But is it really? And if it is (or even if it isn’t), what can we learn from it?

The term ‘enlightenment’ is used in at least two main ways (and probably many others, too). First, there is the sense in which you can talk about the *process* of enlightenment, which can be fast or slow, sudden or gradual. In this sense, there is a path to enlightenment and practices that help people along that path. There can also be temporary experiences of enlightenment, called *kensho* in Zen, and these can be deep or shallow – tiny glimpses or deep experiences of opening.

The neuroscience of meditative practice has gained some traction on these temporary openings-up. Austin (2009) points out that the dramatic shifts of *kensho* are often triggered by unexpected sensory stimuli. So if a practised meditator is deeply absorbed, a sudden stimulus might capture her attention, stop all self-referential default-mode processing, and leave an experience of emptiness without self or time. Austin has no direct evidence that this actually happens in *kensho*, but the possibility hints at a potential coming-together of the neuroscience of self with accounts of the ‘long rigorous path toward selflessness’ (2009, p. 81).

A science of enlightenment experiences is still some way off, not least because the concept of enlightenment is so imprecise compared to the kinds of hypotheses current scientific paradigms are good at testing. This means that there is often disagreement, within and between Buddhist traditions, about which states or traits deserve to be called enlightened, and whether a given individual has attained them. The innovative methods we explored in Chapter 17 might well help here, however. And interestingly, Buddhist teachers tend to practise the same sort of scepticism that scientific inquiry needs: not taking self-reports at face value, but comparing




FIGURE 18.3 • An exercise in headlessness. How many heads are there? Seeing the world this way, you lose your own head and gain everybody else’s.

them against a longer practice history, the manner in which the report is given, and observations of other behaviour (Davis and Vago, 2013). It seems also that meditative practice helps reduce the very biases which this scepticism guards against (and this reduction can be measured using scientifically developed tests), so the prospects seem good for greater dovetailing of scientific and meditative practice.

One of the experiences often thought of in terms of a temporary glimpse of enlightenment is the experience of ‘cessation of all phenomena’, an ‘inseparable emptiness–luminosity–bliss state, not different in nature from awareness itself’ (Davis and Vago, 2013, p. 1). In these moments, there is ‘pure consciousness’: awareness without anything to be aware of or anyone to be aware. Something of this kind has now been studied with two experienced practitioners, comparing these deeper experiences of cessation with the more common experience of a sensory object ‘passing or vanishing from conscious awareness’ (p. 2), and finding much more significant levels of activation at moments of peak clarity in frontal polar cortex – pointing the way to finding potential neural markers relative to individual baseline states. Technological limitations like fMRI’s temporal resolution, and statistical limitations like reliance on a comparison between a state of interest and another state of no interest, will need tackling along with the other questions of methodology we have considered. But for momentary experiences of enlightenment, there seems a clear way forward.

Second, there is what is usually thought of as lying at the end of the path: sometimes called ultimate enlightenment, or full awakening. This is not a state of consciousness like a mystical or religious experience, or even a *kensho* experience, which passes away. Indeed, it is often said not to be a state at all. Everything is just the same as it always was, because everything is inherently enlightened. A famous Zen proverb says, ‘Before enlightenment chop wood, carry water. After enlightenment chop wood, carry water.’ This well describes John Wren-Lewis’s approach to life.

In this sense, there is no path to enlightenment because there is nowhere to go and no one who travels, even though there are many paths that each of us treads. Those who speak of enlightenment at all say that it cannot be explained or described. Anything you say is beside the point. So in this sense



CONCEPT 18.2

PURE CONSCIOUSNESS

Is there such a thing? Pure consciousness is described as a state of wakeful awareness without content of any kind: no thoughts, no perceptions, no self, no other, and no sense of body, space, or time.

This notion appears in Christianity in the medieval work of mysticism *The Cloud of Unknowing*, in Buddhist meditation traditions, and in Hinduism, where *Nirvikalpa Samadhi* is a state of no duality, no mind, and no experience. In transcendental meditation, pure consciousness is said to be what is left when all thoughts cease and the mantra finally drops away. Pure consciousness has been described as a type of mystical experience that cuts right across cultural and linguistic divisions (Forman, 1990, 1999). An attempt has even been made to model this state in a machine that turns its attentional system on itself (Aleksander, 2007, p. 94), and Baars has suggested experimental ways of testing it (in Blackmore, 2005, pp. 20–21).

If you were convinced that pure consciousness exists, this would be a problem for some theories of consciousness that cannot account for it. For example, representationalist theories rely on experience having content, so cannot account for pure consciousness (Bachmann, 2014), and phenomenologists following Husserl claim to have discovered that all experience is intentional, or *about* something. Others have claimed that contentless experience cannot logically exist, that mystical experiences must be shaped by culture and religious training, and that relying on reports of experiences is always dubious (Katz, 1978; Forman, 1990, 1999).

Some neuroscientists are equally dismissive, claiming that ‘it does not make any sense to speak of experience without an *experiencer* who experiences the experiences’ (Cleeremans, 2008, p. 21), let alone to posit experience without content. Hofstadter says that ‘unfortunately for the Zen folks’ we cannot turn off our hallucinations and perceptions. ‘We can try to do so, can tell ourselves we’ve

succeeded, can claim that we have “unperceived” them or whatever, but that’s just self-fooling’ (2007, p. 303). But is it? As Metzinger reminds us, ‘Self-deception may feel like insight’ (2009, p. 220).

This debate has often been confined to arguments over personal experience or theological doctrine, amounting to just another playground scrap: ‘I’ve experienced it, so there!’, ‘Oh no you haven’t’, ‘Oh yes I have’. Yet the possibility or impossibility of pure consciousness may be an example of where the creative self-reflexive methods discussed in Chapter 17 could help.

enlightenment is not a meme, even though the idea of it is. Yet, paradoxically, one person can do things, or point to things, to help others become enlightened, and in this way enlightenment can be passed on. This is known in Zen as ‘transmission outside the scriptures’. This is the point of the koan story about Hui Neng and the monk (Concept 18.1). Perhaps the closest we can get to saying anything positive about enlightenment is that it is losing, not gaining – dropping, or seeing through, the illusions.

All this sounds gloriously paradoxical. It could be glorious nonsense. Or it could be that Zen confronts the same paradoxical problems that the science of consciousness confronts. We have met these many times already. For example, there seems to be both

a private inner world and a public outer world. From this duality arise the hard problem and the explanatory gap. We feel strongly that we, and we alone, know what our inner world is like. Yet as soon as we try to describe it, we find we are providing third-person data, and the special inner world is gone.

We have met illusions too, in perception, and in theories about self and free will. Are these the same illusions that enlightenment sees through? If so, then we might hope to learn something from traditions that have been struggling with the paradoxes and penetrating the illusions for two and a half millennia. If not, then this foray into Buddhism and spirituality will have been a waste of time.

So we return now to our central question. What does all this have to do with a science of consciousness? We saw that both Buddhism and science claim to have




FIGURE 18.4 • Yogins at playtime.

ways of finding out the truth. We can now ask whether it is the same truth that they find.

ILLUSION, NO-SELF, NO DUALITY

From the science and philosophy of consciousness, we have learned that the visual world might be a grand illusion, that the stream of consciousness is not what it seems, and that both self and free will may be illusory too. Buddhist training is aimed at demolishing illusions. So let's look more closely at the Buddhist concept of illusion to see whether it fits with those scientific and philosophical discoveries or not.

The Buddha taught that ordinary experience is illusory because we have wrong, or ignorant, ideas about the nature of the world. We see things, including ourselves, as separately existing entities, when in reality all phenomena are impermanent and empty. This 'emptiness', much spoken of, is not about 'nothingness' or 'voidness'. It means that things are inherently empty of self-nature, or empty of inherent existence. Take a car. This collection of bits and pieces comes together, and for a time we call it 'my car', even if it gets a new engine and replacement exhaust pipes, and then it dissipates into bits again. There is no inherent car-ness there. The illusion is the tendency to treat things as permanent and self-existing. So if someone experiences emptiness during meditation, this does not mean they go into a great void of nothingness; it means that they experience everything that arises as interdependent, impermanent, and not inherently divided into separate things.

This is relatively easy to accept for cars, tables, books, and houses, but much harder when it comes to one's own self. Central to Buddhism is the doctrine of *anatta*, or no-self. Again, this does not mean that the self does not exist (the common English translation is misleading), but that it is conditioned and impermanent like everything else. The Buddha urged people to see things as they are,

to see that what we call 'I', or 'being', is only a combination of physical and mental aggregates, which are working together interdependently in a flux of momentary change within the law of cause and effect, and that there is nothing permanent, everlasting, unchanging and eternal in the whole of existence.

(Rahula, 1959, p. 66)

This is why Derek Parfit (1987) refers to the Buddha as the first bundle theorist.

This theory of no-self went dramatically against the popular beliefs of the Buddha's time, and it goes against the tenets of all the major religions since. Most religions claim that there is a permanent, everlasting entity called a soul or spirit or *atman*. This may survive death to live eternally in heaven or hell, or may go through a series of many lives until it is finally purified and becomes one with God or a universal soul. The Buddha denied all of this and debated the issue with the best thinkers of his time.

Buddhism stands unique in the history of human thought in denying the existence of such a Soul, Self, or *Atman*. According to the teaching of the Buddha, the idea of self is an imaginary, false belief which has no corresponding reality, and it produces harmful thoughts of 'me' and 'mine',

'egolessness or non-self [...] is not an article of faith, but a discovery of mindfulness'

(Mikulas, 2007, p. 32)

'is the self like a unicorn, a mythical being whose representations exist but who is actually imaginary?'

(Hanson and Mendius, 2009, pp. 208–209)

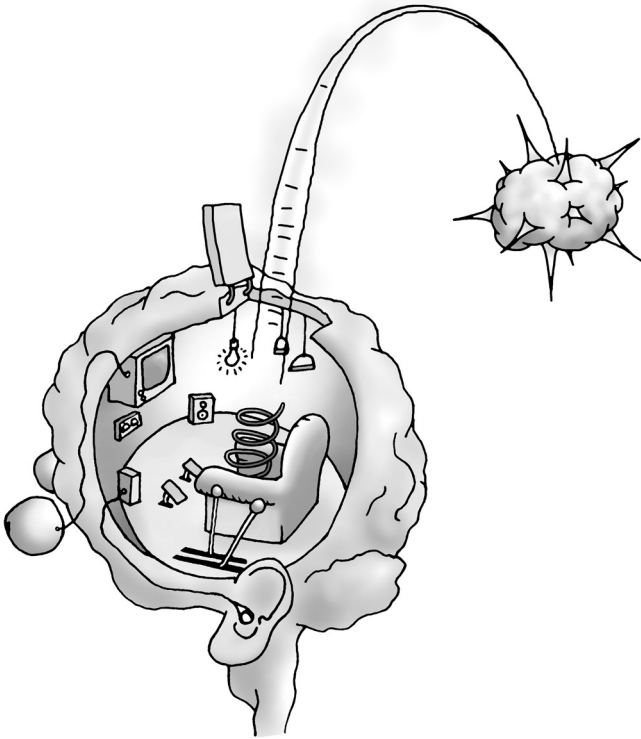


FIGURE 18.5 • Blast-off!

'the fundamental subject/object structure of experience can be transcended'

(Metzinger, 2009, p. 33)

'thought is itself the thinker'

(James, 1890, i, p. 401)

WHAT IS THIS?

'The characteristic of enlightenment is a permanent freeing of the individual from the illusion that he is "doing"'

(Fenwick, 1987, p. 117)

selfish desire, craving, attachment. [. . .] It is the source of all the troubles in the world.

(Rahula, 1959, p. 51)

Even so, it is not easy to give up. A monk once asked the Buddha whether people are ever tormented by finding nothing permanent within themselves. They certainly are, was the response. A man who hears the teachings thinks, 'I will be annihilated, I will be destroyed, I will be no more'. So he weeps, mourns, worries, and becomes bewildered (Rahula, 1959, p. 56). This idea of no-self was just as difficult for people to accept millennia ago as it is now.

Another aspect of the false conception of self is the idea that it can *do* things. The Buddha was clear on this: 'actions do exist, and also their consequences, but the person that acts does not' (Parfit, 1987, p. 21). The Sri Lankan monk and Buddhist scholar Walpola Rahula explains this in words that could have come straight from William James: 'there is no thinker behind the thought. Thought itself is the thinker. If you remove the thought, there is no thinker to be found' (1959, p. 26). Does this mean that free will is an illusion? The question did not arise in early Buddhist cul-

tures and languages in the way it has in the West. Even so, if everything is conditioned and relative, and subject to the law of cause and effect, then it is obvious that nothing can be independent and so truly free (Rahula, 1959). Indeed, the fiction of an independent self that could have freedom is part of the problem, and 'The aim of dharma practice is to free ourselves from this illusion of freedom' (Batchelor, 1997, p. 95).

Relevant here is the Buddhist notion of *karma* or volition. Rahula explains that although the term 'karma' means 'action' or 'doing', in Buddhism it refers only to willed or voluntary actions. These arise from the false idea of a self who thinks and acts, and it is only these kinds of actions that can have consequences that are good or bad. When the false view is dropped, people continue to act, think, and do things, but they no longer accumulate karma because they are free from the false idea that they are a self who acts. Escaping from the wheel of birth and death is, therefore, nothing like the popular idea of reincarnation in which you are rewarded for good actions and punished for bad in a whole series of future lives. Nor is it like being someone who leaves the world of *samsara* and goes to a spiritual realm called *nirvana*. Rather, it means being without the illusion of the self who acts. This is why neuropsychologist Peter Fenwick says that 'The characteristic of enlightenment is a permanent freeing of the individual from the illusion that he is "doing"' (1987, p. 117).

In his classic book *The Way of Zen*, Alan Watts explains how it is just to walk on, wholeheartedly engaged in every action. Yet 'we cannot realize this kind of action until it is clear beyond any shadow of doubt that it is actually impossible

to do anything else' (1957, p. 161). This is 'unmotivated non-volitional functioning'. It is 'non-action' or 'not-doing'. It is how things are because really there is no entity to act; no entity to be either bound or free (Wei Wu Wei, 2004). Wu Wei (whose name comes from Taoism) suggests 'asking yourself whether you are not still looking as from a phenomenal centre that has only an imaginary existence. If so, you will be misled; if not – you will understand at once' (Watts, 1957, p. 163).



PRACTICE 18.2

MINDFULNESS

Your final task is to be mindful for a whole day (or forever, if you prefer). If possible, choose a day when you will have time on your own, and when you might be walking, doing housework, gardening, or taking part in sports, rather than reading, writing, and socialising. Decide that you will stay fully present in every moment and then begin. You must begin – and continue – with *this* moment and not think about how well you have done so far, or how long you still have to go. Just attend, fully and clearly, to what is going on *now*. You will probably find that it seems easy to begin with, and that everything seems bright and clear when you do, but then you will suddenly realise that you have gone off into some train of thought and lost the mindfulness. Do not get cross with yourself but just return to the present moment. That's all you have to do.

It is very difficult. Don't get discouraged.

You might like to make notes on how you got on, or discuss the following questions later with friends. What made it harder or easier to maintain mindfulness? Were you ever frightened? Did being mindful interfere with what you were doing? How does this task relate to all the previous ones? Can you imagine being mindful all your life?

What is it like being mindful?

'I do the "as if". And I think almost everybody who's happy and healthy tends to do that.'

(Dan Wegner, in Blackmore, 2005, p. 257)

How is it possible to live without doing? One way lies in the simple phrase 'as if'. You can live *as if* you have free will; *as if* you are a self who acts; *as if* there is a physical world outside yourself. You can treat others *as if* they are sentient beings who have desires, beliefs, hopes, and fears. In discussing free will with scientists and philosophers (Blackmore, 2005), Sue discovered that this compromise is a common solution. Alternatively, you can throw off the idea altogether and simply accept that all your own actions are just part of the inevitable play of the whole amazing, complex universe of which your ever-renewing body and illusory self are part. This way of living drops any distinction between real and *as-if* intentionality, or real and *as-if* free will, and drops the illusion of a self who acts (Blackmore, 2013).

Does any of this help us with the hard problem, and with the dualism that bedevils every attempt to make scientific sense of consciousness? It is said that when

people drop all the illusions, nonduality is revealed, and 'there is no longer any vestige of a distinction between self and experience' (Claxton, 1986b, p. 319). In Buddhism, this is likened to polishing a mirror. When the mirror is completely spotless, there is no distinction between the world and its reflection, and the mirror disappears.

WHAT IS IT LIKE BEING MINDFUL?

Have these people really seen nonduality, directly, in their own experience? If they have, could we all see it? Might the psychologists, philosophers, neuroscientists, and all other thinkers working on the problem of consciousness see nonduality directly for themselves? If so, it seems possible that they might bring together two apparently totally different sets of disciplines: the disciplines of science and the disciplines of self-transformation. They might then understand exactly what had happened in their own brains and the rest of their bodies when all the illusions fell away and the distinctions between the first, the second, and the third person were gone. This way the direct experience of nonduality might be integrated into a study of consciousness which at the moment knows intellectually that dualism must be false, but has not yet worked out what this really means, or how to take the next step.

'experiencing is impossible without an experiencer'

(Strawson, 2011, p. 254)

The prospects for a profound integration of this kind seem real, when we consider all the parallels between what we have learned in this book and what the practice of nonduality teaches us. We have explored the possibility that visual experience and perhaps consciousness as a whole may be subject to a 'grand illusion'. We have surveyed findings which challenge the intuition that action or attention happens because 'I' make a decision to act or attend. We have contemplated the idea that it is impossible to find a function for consciousness beyond the interactions of everything else in our bodies and environments. We have asked whether it is impossible to pin down hard boundaries between ordinary and altered states, or between what is real and what is imagined, or between animals and machines that are or aren't conscious. We have gathered a vast amount of evidence from different corners of consciousness studies, much of which turns out to support the basic spiritual notions of impermanence of self and continuity of self with the rest of the universe.

'This is just the universeing'

(write in your own name)

So it may be that the deepest mystical insights are not only monist and non-paranormal, but are perfectly compatible with the world described by physics (Hunt, 2006). Perhaps the experience of unity or oneness that is so common in mystical and psychedelic experiences is a valid insight into an ultimately unified and integrated universe in which everything affects everything else. This might be summarised as 'the universe is one, the separate self is an illusion, immortality is not in the future but now, and there is nothing to be done'.

If these insights are valid, what needs overthrowing is not monist science but the vestiges of dualist thinking that still lurk within it. This idea gives no comfort to those who hope for personal survival of death, but it is compatible with our scientific understanding of the universe.

We have also begun to learn why this is often not how it feels: why we fall for free will, pictures in the head, and singular streams of consciousness, even if they are fictions; why we like the clean lines between conscious and unconscious, voluntary and involuntary, human and machine, self and other, inner

and outer, even if we know we have invented them. There are reasons why all these ways of thinking come about, why there are so many commonalities across individuals and societies, and also why there are some individual and cultural variations. Illusions should not just be dismissed as stupid mistakes. They are ways of seeing the world as other than it is – and that world includes your own experience. There are always reasons for these other ways of seeing. Maybe we have the illusions we have because they are a handy simplification of reality. Maybe we have them because they help us feel a sense of control in an uncontrollable universe. But if we really want to answer the question of how and why we have conscious experience (or think we do), we need to question those reasons.

If the conclusion we draw from the evidence we have considered in this book is that what needs solving is not the hard problem but the problem of why we feel there is a hard problem, then the questions to be asked change, too. Instead of banging our heads against the brick wall that separates the activity of the nervous system from *the experience itself*, we can turn to a new set of questions.

How do our embodied cognitive capacities give rise to the illusion of consciousness?

Is the illusion of consciousness evolutionarily adaptive (or was it once upon a time but no longer)?

How do we learn to let go of the illusion?

And what happens when we do?

Whatever questions we are asking, we should always remember to preserve a healthy scepticism about our own experiences – even ones that feel like profound awakenings. In his book *Waking, Dreaming, Being* (2014), Evan Thompson talks about waking and dreaming and the Indian myth of the receding frame. He reminds us that we can confirm that we're dreaming by waking up – either by waking from a dream or by becoming lucid within a dream. But we can never confirm that we're awake because there's always a chance that we might really wake up. 'The reason is that for any experience we choose – specifically, any experience we take to be a waking one – it seems conceivable that we could wake up from that experience' (2014, p. 194). It is a valuable – and an interesting – practice to ask yourself now and then: do I think I am awake now, or do I think I have just woken up, and if so why, and might I be mistaken?

If we turn our attention to this new set of questions, and strip away as many assumptions as we can about our own wakefulness, will the problem of consciousness be solved? We do not know. Zen is said to require 'great doubt', great determination, and the more perplexity the better. The same might be said of a science of consciousness. We hope that you, like us, are now more perplexed than when you began – but with a little glimpse, too, of what might lie beyond.

'We need to say 'I don't know' often [so] we do not fall entirely under the enchanted spell of our own standpoint'

(Saunders, 2014, p. 187)

'motivation for Zen training lies in perplexity'

(Crook, 1990, p. 171)

READING

Blackmore, S. (2011). *Zen and the art of consciousness*. London: Oneworld.

A personal view of meditation is described in the Introduction (pp. 4–15). Try any of the questions, perhaps especially ‘What am I doing?’ (pp. 135–149). Students could be assigned, or choose, one chapter to read and present in class. Regular meditators might like to work with one of these koans and report on their experiences.

Claxton, G. (1996a). The light’s on but there’s nobody home: The psychology of no-self. In G. Claxton (Ed.), *Beyond therapy: The impact of Eastern religions on psychological theory and practice* (pp. 49–70). Dorset: Prism.

What happens when you ‘give up the ghost’?

Metzinger, T. (2009). Consciousness technologies and the image of humankind. In T. Metzinger, *The ego tunnel* (pp. 207–218). New York: Basic Books.

The final phase of the Consciousness Revolution: the science, philosophy, and ethics of what a human being (or an ego machine) is and what it should become.

Mikulas, W.L. (2007). Buddhism & Western psychology: Fundamentals of integration. *Journal of Consciousness Studies*, 14(4), 4–49.

Describes ‘essential Buddhism’ as psychology, not religion or philosophy, arguing for its integration with mainstream ‘Western’ psychology to better understand the mind and its disorders.

Rosch, E. (1997). Transformation of the Wolf Man. In J. Pickering (Ed.), *The authority of experience: Essays on Buddhism and psychology* (pp. 6–27). Richmond, Surrey: Curzon.

A useful basis for discussing the question 'Are psychotherapy and spiritual development the same or different?'





FIGURE 18.6

'I don't think –' 'Then you shouldn't talk,' said the Hatter.

(Lewis Carroll, Alice's Adventures in Wonderland, 1865)



References

- Aaronson, S.** (2014). Why I am not an Integrated Information Theorist (or, The unconscious expander). *Shtetl-Optimized*. www.scottaaronson.com/blog/?p=1799
- Adams, D.** (1979). *The hitch hiker's guide to the galaxy*. London: Pan.
- Adams, S. S., and Burbeck, S.** (2012). Beyond the octopus: From general intelligence toward a human-like mind. In P. Wang and B. Goertzel (Eds), *Theoretical foundations of artificial general intelligence* (pp. 49–65). Amsterdam: Atlantis.
- Adolphs, R.** (2015). The unsolved problems of neuroscience. *Trends in Cognitive Sciences*, 19(4), 173–175.
- Afraz, S.-R., Kiani, R., and Esteky, H.** (2006). Microstimulation of inferotemporal cortex influences face categorization. *Nature*, 442, 692–695.
- Aglioti, S., Goodale, M.A., and DeSouza, J.F.X.** (1995). Size contrast illusions deceive the eye but not the hand. *Current Biology*, 5, 679–685.
- Akins, K.A.** (1993). What is it like to be boring and myopic? In B. Dahlbom (Ed.), *Dennett and his critics: Demystifying mind* (pp. 124–160). Oxford: Blackwell.
- Alais, D., Cass, J., O'Shea, R.O., and Blake, R.** (2010). Visual sensitivity underlying changes in visual consciousness. *Current Biology*, 20(15), 1362–1367.
- Albahari, M.** (2006). *Analytical Buddhism: The two-tiered illusion of self*. New York: Palgrave Macmillan.
- Albrecht, T., and Mattler, U.** (2012). Individual differences in subjective experience and objective performance in metacontrast masking. *Journal of Vision*, 12(5), 5, 1–24.
- Alderson-Day, B., and Fernyhough, C.** (2016). Auditory verbal hallucinations: Social, but how? *Journal of Consciousness Studies*, 23(7–8), 163–194.
- Aleksander, I.** (2005). *The world in my mind, my mind in the world*. Exeter: Imprint Academic.

• R E F E R E N C E S

Aleksander, I. (2007). Machine consciousness. In M. Velmans and S. Schneider (Eds), *The Blackwell companion to consciousness* (pp. 87–98). Oxford: Blackwell.

Aleksander, I., and Morton, H. (2007). Why axiomatic models of being conscious? *Journal of Consciousness Studies*, 14, 15–27.

Alkire, M.T., Haier, R.J. and Fallon, J.H. (1998). Toward the neurobiology of consciousness: Using brain imaging and anesthesia to investigate the anatomy of consciousness. In S. R. Hameroff, A. W. Kaszniak, and A. C. Scott (Eds), *Toward a science of consciousness II: The Second Tucson Discussions and Debates* (pp. 255–268). Cambridge, MA: MIT Press.

Alkire, M., Hudetz, A.G., and Tononi, G. (2008). Consciousness and anesthesia. *Science*, 322, 876–880.

Alkire, M.T., and Miller, J. (2005). General anaesthesia and the neural correlates of consciousness. *Progress in Brain Research*, 150, 229–250.

Allen, C., and Trestman, M. (2016). Animal consciousness. In S. Schneider and M. Velmans (Eds), *The Blackwell companion to consciousness* (pp. 63–76). 2nd ed. Chichester, West Sussex: Wiley Blackwell.

Allport, A. (1993). Attention and control: Have we been asking the wrong questions? A critical review of twenty-five years. In D. E. Myer and S. Kornblum (Eds), *Attention and Performance XIV* (pp. 183–218). Cambridge, MA: MIT Press.

Allport, A. (2011). Attention and integration. In C. Mole, D. Smithies, and W. Wu (Eds), *Attention: Philosophical and psychological essays* (pp. 24–59). New York: Oxford University Press.

Allred, S., Anderson, B., Brainard, D.H., Gegenfurtner, K., and Maloney, L.T. (Eds). (2017, January). A dress rehearsal for vision science. *Journal of Vision*, 17(1). <http://jov.arvojournals.org/SS/thedress.aspx>

Alpert, R. (Baba Ram Dass) (1971). *Be here now*. San Cristobal, NM: Lama Foundation.

Alsmith, A.J., and Longo, M. R. (2014). Where exactly am I? Self-location judgements distribute between head and torso. *Consciousness and Cognition*, 24, 70–74.

Alter, T. (2010). A defense of the necessary unity of phenomenal consciousness. *Pacific Philosophical Quarterly*, 91, 19–37.

Alvarado, C.S. (1982). ESP during out-of-body experiences: A review of experimental studies. *Journal of Parapsychology*, 46, 209–230.

Alvarado, C.S., and Zingrone, N.L. (2015). Features of out-of-body vexpériences: Relationships to frequency, wilfulness and previous knowledge about the experience. *Journal of the Society for Psychological Research*, 79, 98–111.

- American Psychiatric Association.** (2013). *Diagnostic and statistical manual of mental disorders (DSM-5®)*. American Psychiatric Pub.
- Anderson, B.** (2011). There is no such thing as attention. *Frontiers in Psychology*, 2, article 246.
- Andrade, J.** (2012). Consciousness. In N. Braisby and A. Gellatly (Eds), *Cognitive psychology* (pp. 577–605). 2nd ed. Oxford: Oxford University Press.
- Anon.** (14th century/2009). *The cloud of unknowing. With the Book of privy counsel*. Trans. C. A. Butcher. Boston, MA: Shambhala.
- Aru, J., Bachmann, T., Singer, W., and Melloni, L.** (2012). Distilling the neural correlates of consciousness. *Neuroscience & Biobehavioral Reviews*, 36(2), 737–746.
- Aru, J., and Bachmann, T.** (2015). Still wanted – the mechanisms of consciousness! *Frontiers in Psychology*, 6, article 5.
- Arzy, S., Seeck, M., Ortigue, S., Spinelli, L., and Blanke, O.** (2006). Induction of an illusory shadow person. *Nature*, 443, 287.
- Asch, S.E.** (1952). *Social psychology*. Englewood Cliffs, NJ: Prentice-Hall.
- Assefa, S.Z., Diaz-Abad, M., Wickwire, E.M., and Scharf, S.M.** (2015). The functions of sleep. *AIMS Neuroscience*, 2(3), 155–171.
- Aunger, R.A.** (Ed.) (2000). *Darwinizing culture: The status of memetics as a science*. Oxford: Oxford University Press.
- Austen, J.** (1813). *Pride and prejudice*. London: T. Egerton. Full text available at <https://www.gutenberg.org/files/1342/1342-h/1342-h.htm> and <https://books.google.co.uk/books?id=vKYBAAAQAAJ>
- Austin, J.H.** (1998). *Zen and the brain: Toward an understanding of meditation and consciousness*. Cambridge, MA: MIT Press.
- Austin, J. H.** (2006). *Zen-brain reflections: Reviewing recent developments in meditation and states of consciousness*. Cambridge, MA: MIT Press.
- Austin, J.H.** (2009). *Selfless insight: Zen and the meditative transformations of consciousness*. Cambridge, MA: MIT Press.
- Awh, E., Belopolsky, A.V., and Theeuwes, J.** (2012). Top-down versus bottom-up attentional control: A failed theoretical dichotomy. *Trends in Cognitive Sciences*, 16(8), 437–443.
- Azzi, J.C., Gattass, R., Lima, B., Soares, J.G.M., and Fioerani, M.** (2015). Precise visuotopic organization of the blind spot representation in primate V1. *Journal of Neurophysiology*, 113, 3588–3599.
- Baars, B.J.** (1988). *A cognitive theory of consciousness*. Cambridge: Cambridge University Press.
- Baars, B.J.** (1997a). In the theatre of consciousness: Global workspace theory, a rigorous scientific theory of consciousness. *Journal of Consciousness Studies*, 4, 292–309. Commentaries and author's response pp. 310–364.

• R E F E R E N C E S

- Baars, B.J.** (1997b). *In the theater of consciousness: The workspace of the mind*. New York: Oxford University Press.
- Baars, B.J.** (1999). There is already a field of systematic phenomenology, and it's called 'psychology'. *Journal of Consciousness Studies*, 6(2–3), 216–218. Also in F.J. Varela and J. Shear (Eds) (1999). *The view from within* (pp. 216–218). Thorverton, Devon: Imprint Academic.
- Baars, B.J.** (2005a). Global workspace theory of consciousness: Toward a cognitive neuroscience of human experience. *Progress in Brain Research*, 150, 45–53.
- Baars, B.J.** (2005b). Subjective experience is probably not limited to humans: The evidence from neurobiology and behaviour. *Consciousness and Cognition*, 14, 7–21.
- Baars, B.** (2012). The biological cost of consciousness. *Nature Precedings*. <http://precedings.nature.com/documents/6775/version/1>
- Baars, B.J., and Franklin, S.** (2009). Consciousness is computational: The LIDA model of global workspace theory. *International Journal of Machine Consciousness*, 1, 23–32.
- Baars, B.J., and Gage, N.M.** (2010). *Cognition, brain and consciousness: Introduction to cognitive neuroscience*. 2nd ed. Burlington, MA: Academic.
- Bachmann, J.K.** (2014). *Accounting for pure consciousness: An examination of the ability of the representationalist approach to phenomenal consciousness to account for pure consciousness experiences*. PhD thesis, University of Alberta.
- Bachmann, T., and Hudetz, A.G.** (2014). It is time to combine the two main traditions in the research on the neural correlates of consciousness: C = L x D. *Frontiers in Psychology*, 5, article 940.
- Bach-y-Rita, P.** (1995). *Nonsynaptic diffusion, neurotransmission and late brain reorganization*. New York: Demos.
- Bach-y-Rita, P., and González, J.C.** (2002). *Tactile sensory substitution in blind subjects*. Paper presented at Toward a Science of Consciousness, Tucson, AZ April 2002. Conference Research Abstracts (provided by *Journal of Consciousness Studies*), Abstract No. 186.
- Baddeley, A.D.** (2000). Short-term and working memory. In E. Tulving and F.I.M. Craik (Eds), *The Oxford handbook of memory* (pp. 77–92). New York: Oxford University Press.
- Baer, J., Kaufman, J.C., and Baumeister, R.F.** (2008). *Are we free? Psychology and free will*. New York: Oxford University Press.
- Bagchi, B.K., and Wenger, M.** (1957). Electrophysiological correlates of some yogic exercises. *Electroencephalography and Clinical Neurophysiology*, 10, 132–149.
- Bahrack, L.E., Moss, L., and Fadil, C.** (1996). Development of visual self-recognition in infancy. *Ecological Psychology*, 8(3), 189–208.

- Baker, D.H.** (2010). Visual consciousness: The binocular rivalry explosion. *Current Biology*, 20(15), R644–646.
- Balcombe, J.** (2016). *What a fish knows: The inner lives of our underwater cousins*. New York: Scientific American/Farrar, Straus and Giroux.
- Banks, W.P.** (1993). Problems in the scientific pursuit of consciousness. *Consciousness and Cognition*, 2, 255–263.
- Barbey, A.K.** (2017). Network neuroscience theory of human intelligence. *Trends in Cognitive Sciences*. [http://www.cell.com/trends/cognitive-sciences/fulltext/S1364-6613\(17\)30221-8](http://www.cell.com/trends/cognitive-sciences/fulltext/S1364-6613(17)30221-8)
- Barbur, J.L., Watson, J.D.G., Frackowiak, R.S.J., and Zeki, S.** (1993). Conscious visual perception without V1. *Brain*, 116, 1293–1302.
- Barkow, J.H., Cosmides, L., and Tooby, J.** (Eds) (1992). *The adapted mind: Evolutionary psychology and the generation of culture*. Oxford: Oxford University Press.
- Barlow, H.** (1987). The biological role of consciousness. In C. Blakemore and S. Greenfield (Eds), *Mindwaves* (pp. 361–374). Oxford: Blackwell.
- Baron-Cohen, S., and Harrison, J.** (Eds) (1997). *Synaesthesia: Classic and contemporary readings*. Oxford: Blackwell.
- Baron-Cohen, S., Johnson, D., Asher, J., Wheelwright, S., Fisher, S. E., Gregersen, P. K., and Allison, C.** (2013). Is synaesthesia more common in autism? *Molecular Autism*, 4(1), 40.
- Barrett, D., and McNamara, P.** (2012). *Encyclopedia of sleep and dreams: The evolution, function, nature, and mysteries of slumber*. 2 vols. Santa Barbara, CA: Greenwood.
- Batchelor, M.** (2001). *Meditation for life*. London: Frances Lincoln.
- Batchelor, S.** (1990). *The faith to doubt: Glimpses of Buddhist uncertainty*. Berkeley, CA: Parallax Press.
- Batchelor, S.** (1994). *The awakening of the West: The encounter of Buddhism and Western culture*. London: Aquarian.
- Batchelor, S.** (1997). *Buddhism without beliefs: A contemporary guide to awakening*. London: Bloomsbury.
- Batchelor, S.** (2015). *After Buddhism: Rethinking the dharma for a secular age*. New Haven, CT: Yale University Press.
- Bauby, J.-D.** (1997). *The diving bell and the butterfly [Le Scaphandre et le papillon]*. Trans. J. Leggatt. London: Fourth Estate.
- Baumeister, R. F., Masicampo, E. J., and DeWall, C. N.** (2009). Prosocial benefits of feeling free: Disbelief in free will increases aggression and reduces helpfulness. *Personality and Social Psychology Bulletin*, 35, 260–268.

• R E F E R E N C E S

- Baxter, S.** (1999/2015). *Manifold: Time*. London: HarperCollins.
- Bayne, T.** (2005). Divided brains and unified phenomenology: A review essay on Michael Tye's *Consciousness and persons*. *Philosophical Psychology*, 18(4), 495–512.
- Bayne, T.** (2011). Libet and the case for free will skepticism. In R. Swinburne (Ed.), *Free will and modern science* (pp. 25–46). Oxford: Oxford University Press.
- Bayne, T., and Chalmers, D.** (2003). What is the unity of consciousness? In A. Cleeremans (Ed.), *The unity of consciousness: Binding, integration and dissociation* (pp. 23–58). New York: Oxford University Press.
- Bayne, T., Cleeremans, A., and Wilken, P.** (2009). *The Oxford companion to consciousness*. Oxford: Oxford University Press.
- Beaton, M.** (2005). What RoboDennett still doesn't know. *Journal of Consciousness Studies*, 12, 3–25.
- Beierholm, U. R., Quartz, S. R., and Shams, L.** (2009). Bayesian priors are encoded independently from likelihoods in human multisensory cortex. *Journal of Vision*, 9(5), 23, 1–9.
- Bem, D. J.** (2011). Feeling the future: Experimental evidence for anomalous retroactive influences on cognition and affect. *Journal of Personality and Social Psychology*, 100(3), 407–425.
- Bem, D. J., and Honorton, C.** (1994). Does psi exist? Replicable evidence for an anomalous process of information transfer. *Psychological Bulletin*, 115, 4–18.
- Bem, D. J., Palmer, J., and Broughton, R. S.** (2001). Updating the ganzfeld database: A victim of its own success? *Journal of Parapsychology*, 65, 207–218.
- Bennett, C. M., Miller, M. B., and Wolford, G. L.** (2009). Neural correlates of interspecies perspective taking in the post-mortem Atlantic Salmon: An argument for multiple comparisons correction. *NeuroImage*, 47(Suppl. 1), S125.
- Bennett, M. R., and Hacker, P.M.S.** (2003). *Philosophical foundations of neuroscience*. Oxford: Blackwell.
- Berger, J.** (1972). *Ways of seeing*. London: Penguin.
- Bergson, H.** (1920). *Mind-energy: Lectures and essays [L'Énergie spirituelle. Essais et conférences]*. Trans. H. W. Carr. New York: Henry Holt.
- Bering, J. M.** (2002). Intuitive conceptions of dead agents' minds: The natural foundations of afterlife beliefs as phenomenological boundary. *Journal of Cognition and Culture*, 2, 263–308.
- Bering, J. M., and Bjorklund, D. F.** (2004). The natural emergence of reasoning about the afterlife as a developmental regularity. *Developmental Psychology*, 40, 217–233.

- Bertossa, F., Besa, M., Ferrari, R., and Ferri, F.** (2008). Point zero: A phenomenological inquiry into the seat of consciousness. *Perceptual and Motor Skills*, 107(2), 323–335.
- Beth, T., and Ekroll, V.** (2015). The curious influence of timing on the magic experience evoked by conjuring tricks involving false transfer: Decay of amodal object permanence? *Psychological Research*, 79, 513–522.
- Billig, M.** (2012). Abraham Tucker as an 18th-century William James: Stream of consciousness, role of examples, and the importance of writing. *Theory & Psychology*, 22(1), 114–129.
- Bingham, G. P., Coats, R., and Mon-Williams, M.** (2007). Natural prehension in trials without haptic feedback but only when calibration is allowed. *Neuropsychologia*, 45, 288–294.
- Birn, R. M., Diamond, J. B., Smith, M. A., and Bandettini, P. A.** (2006). Separating respiratory-variation-related fluctuations from neuronal-activity-related fluctuations in fMRI. *NeuroImage*, 31(4), 1536–1548.
- Bisenius, S., Trapp, S., Neumann, J., and Schroeter, M. L.** (2015). Identifying neural correlates of visual consciousness with ALE meta-analyses. *NeuroImage*, 122, 177–187.
- Bisiach, E.** (1988). The (haunted) brain and consciousness. In A. J. Marcel and E. Bisiach (Eds), *Consciousness in contemporary science* (pp. 101–120). Oxford: Oxford University Press.
- Bisiach, E.** (1992). Understanding consciousness: Clues from unilateral neglect and related disorders. In A. D. Milner and M. D. Rugg (Eds), *The neuropsychology of consciousness* (pp. 113–137). London: Academic Press. Reprinted in N. Block, O. Flanagan, and G. Güzeldere (Eds) (1991), *The nature of consciousness: Philosophical debates* (pp. 237–253). Cambridge, MA: MIT Press.
- Bisiach, E., and Luzzatti, C.** (1978). Unilateral neglect of representational space. *Cortex*, 14, 129–133.
- Blackmore, S. J.** (1982). *Beyond the body: An investigation of out-of-the-body experiences*. London: Heinemann. Reprinted (1992) with new postscript, Chicago: Academy Chicago.
- Blackmore, S. J.** (1987). A report of a visit to Carl Sargent's laboratory. *Journal of the Society for Psychical Research*, 54, 186–198.
- Blackmore, S. J.** (1992). Psychic experiences: Psychic illusions. *Skeptical Inquirer*, 16, 367–376.
- Blackmore, S. J.** (1993). *Dying to live: Science and the near-death experience*. London: Grafton.
- Blackmore, S. J.** (1996a). *In search of the light: The adventures of a parapsychologist*. Amherst, NY: Prometheus.
- Blackmore, S. J.** (1996b). Out-of-body experiences. In G. Stein (Ed.), *Encyclopedia of the paranormal* (pp. 471–483). Buffalo, NY: Prometheus.

• R E F E R E N C E S

Blackmore, S. J. (1997). Probability misjudgment and belief in the paranormal: A newspaper survey. *British Journal of Psychology*, *88*, 683–689.

Blackmore, S. J. (1998). Why psi tells us nothing about consciousness. In S. R. Hameroff, A. W. Kaszniak, and C. Scott (Eds), *Toward a science of consciousness II* (pp. 710–707). Cambridge, MA: MIT Press.

Blackmore, S. J. (1999). *The meme machine*. Oxford: Oxford University Press.

Blackmore, S. (2001). Three experiments to test the sensorimotor theory of vision. Commentary on O'Regan and Noë, *Behavioral and Brain Sciences*, *24*(5), 977.

Blackmore, S. J. (2002). There is no stream of consciousness. *Journal of Consciousness Studies*, *9*(5–6), 17–28.

Blackmore, S. (2004). *A retroselection theory of dreams*. Poster presented at the Association for the Scientific Study of Consciousness, ASSC8 Antwerp, Belgium, 25–28 June 2004. www.susanblackmore.co.uk/conferences/a-retroselection-theory-of-dreams/

Blackmore, S. (2005). *Conversations on consciousness: What the best minds think about the brain, free will, and what it means to be human*. Oxford: Oxford University Press

Blackmore, S. (2007a). Seeing or blind? A test of sensorimotor theory. Poster presented at the conference 'Perception, action and consciousness: Sensorimotor dynamics and dual vision', Bristol, 1–3 July.

Blackmore, S. (2007b). Memes, minds and imagination. In I. Roth (Ed.), *Imaginative minds* (pp. 61–78). Oxford: Oxford University Press.

Blackmore, S. (2009). A psychological theory of the OBE. In C. D. Murray (Ed.), *Psychological scientific perspectives on out of body and near death experiences* (pp. 23–36). New York: Nova. Reprinted (with new postscript) from (1984), *Journal of Parapsychology*, *48*, 201–218.

Blackmore, S. (2010). Memetics does provide a useful way of understanding cultural evolution. In F. Ayala and R. Arp (Eds), *Contemporary debates in philosophy of biology* (pp. 255–272). Chichester: Wiley-Blackwell.

Blackmore, S. (2011). *Zen and the art of consciousness*. London: Oneworld. Originally published (2009) as *Ten Zen Questions*.

Blackmore, S. (2012). Turning on the light to see how the darkness looks. In S. Kreidler and O. Maimon (Eds), *Consciousness: Its nature and functions*. New York: Nova. www.susanblackmore.co.uk/chapters/turning-on-the-light-to-see-how-the-darkness-looks/

Blackmore, S. (2013). Living without free will. In G. D. Caruso (Ed.), *Exploring the illusion of free will and moral responsibility* (pp. 161–175). Lanham, MD: Lexington Books.

- Blackmore, S.** (2014). Are you convinced that dreaming is a conscious state? In N. Tranquillo (Ed.), *Dream consciousness: Allan Hobson's new approach to the brain and its mind* (pp. 91–93). Cham: Springer International.
- Blackmore, S. J.** (2016). Delusions of consciousness. *Journal of Consciousness Studies*, 23(11–12), 52–64.
- Blackmore, S.** (2016). *Jinny Jana's giant journeys*. Unpublished manuscript.
- Blackmore, S.** (2017). *Seeing myself: The new science of out-of-body experiences*. London: Robinson.
- Blackmore, S. J., Brelstaff, G., Nelson, K., and Troscianko, T.** (1995). Is the richness of our visual world an illusion? Transsaccadic memory for complex scenes. *Perception*, 24, 1075–1081.
- Blackmore, S. J., and Hart-Davis, A.** (1995). *Test your psychic powers*. London: Thorsons. Also (1997) New York: Sterling. (Also in Kindle edition.)
- Blagrove, M.** (2009). Dreaming, scientific perspectives. In T. Bayne, A. Cleeremans, and P. Wilken (Eds), *The Oxford companion to consciousness* (pp. 240–243). Oxford: Oxford University Press.
- Blanke, O., and Arzy, S.** (2005). The out-of-body experience: Disturbed self-processing at the temporo-parietal junction. *Neuroscientist*, 11, 16–24.
- Blanke, O., and Mohr, C.** (2005). Out-of-body experience, heautoscopy, and autoscopic hallucination of neurological origin: Implications for neuro-cognitive mechanisms of corporeal awareness and self-consciousness. *Brain Research Reviews*, 50(1), 184–199.
- Blanke, O., Mohr, C., Michel, C. M., Pascual-Leone, A., Brugger, P., Seeck, M., Landis, T., and Thut, G.** (2005). Linking out-of-body experience and self processing to mental own-body imagery at the temporoparietal junction. *Journal of Neuroscience*, 25, 550–557.
- Blanke, O., Ortigue, S., Landis, T., and Seeck, M.** (2002). Stimulating illusory own-body perceptions. *Nature*, 419, 269–270.
- Blanke, O., Ortigue, S., Spinelli, L., and Seeck, M.** (2004). Out-of-body experience and autoscopia of neurological origin. *Brain*, 127, 243–258.
- Block, N.** (1995). On a confusion about a function of consciousness. *Behavioral and Brain Sciences*, 18(2), 227–287 (incl. commentaries and author's response).
- Block, N.** (2005). Two neural correlates of consciousness. *Trends in Cognitive Sciences*, 9(2), 46–52.
- Block, N.** (2007). Consciousness, accessibility, and the mesh between psychology and neuroscience. *Behavioral and Brain Sciences*, 30(5–6), 481–548 (incl. commentaries and author's response).
- Block, N.** (2010). Attention and mental paint. *Philosophical Issues*, 20, 23–63.

• R E F E R E N C E S

Block, N. (2011). Perceptual consciousness overflows cognitive access. *Trends in Cognitive Sciences*, 15(12), 567–575.

Bloom, P. (2004). *Descartes' baby: How child development explains what makes us human*. London: Heinemann.

Boly, M., Seth, A. K., Wilke, M., Ingmundson, P., Baars, B., Laureys, S., Edelmans, D. B., and Tsuchiya, N. (2013). Consciousness in humans and non-human animals: Recent advances and future directions. *Frontiers in Psychology*, 4, article 625.

Booth, M. (2003). *Cannabis: A history*. London: Doubleday.

Boswell, J. (1791/1952). *Life of Samuel Johnson, LL.D. Great Books of the Western World*, vol. 44. Chicago, IL: W. Benton: Encyclopædia Britannica.

Botvinick, M., and Cohen, J. (1998). Rubber hands 'feel' touch that eyes see. *Nature*, 391, 756.

Bouso, J. C., Doblin, R., Farre, M., Alcazar, M. A., and Gomez-Jarabo, G. (2008). MDMA-assisted psychotherapy using low doses in a small sample of women with chronic posttraumatic stress disorder. *Journal of Psychoactive Drugs*, 40, 225–236.

Boyd, R., and Richerson, P. J. (2009). Culture and the evolution of human cooperation. *Philosophical Transactions of the Royal Society B*, 364(1533), 3281–3288.

Braithwaite, J. J. (2008). Towards a cognitive neuroscience of the dying brain. *Skeptical*, 21, 8–16.

Braithwaite, J. J., James, K., Dewe, H., Medford, N., Takahashi, C., and Kessler, K. (2013). Fractionating the unitary notion of dissociation: Disembodied but not embodied dissociative experiences are associated with exocentric perspective-taking. *Frontiers in Human Neuroscience*, 7, article 719.

Brambilla, M., Ferrante, E., Birattari, M., and Dorigo, M. (2013). Swarm robotics: A review from the swarm engineering perspective. *Swarm Intelligence*, 7(1), 1–41.

Brandl, J. L. (2016). The puzzle of mirror self-recognition. *Phenomenology and Cognitive Science*. Online first.

Brasington, L. (2015). *Right concentration: A practical guide to the jhānas*. Boston, MA: Shambhala.

Brass, M., Lynn, M. T., Demanet, J., and Rigoni, D. (2013). Imaging volition: what the brain can tell us about the will. *Experimental Brain Research*, 229(3), 301–312.

Breazeal, C. L. (2001). *Designing sociable robots*. Cambridge, MA: MIT Press.

Breitmeyer, B. G. (2015). Psychophysical 'blinding' methods reveal a functional hierarchy of unconscious visual processing. *Consciousness and Cognition*, 35, 234–250.

Bressloff, P. C., Cowan, J. D., Golubitsky, M., Thomas, P. J., and Wiener, M. C. (2002). What geometric visual hallucinations tell us about the visual cortex. *Neural Computation*, 14, 473–491.

Brewer, J. A., Worhunksy, P. D., Gray, J. R., Tang, Y. Y., Weber, J., and Kober, H. (2011). Meditation experience is associated with differences in default mode network activity and connectivity. *Proceedings of the National Academy of Sciences, USA*, 108(50), 20254–20259.

Bridgeman, B., Lewis, S., Heit, G., and Nagle, M. (1979). Relation between cognitive and motor-oriented systems of visual position perception. *Journal of Experimental Psychology: Human Perception and Performance*, 5, 692–700.

Broadbent, D. E. (1958). *Perception and communication*. New York: Pergamon Press.

Broks, P. (2003). *Into the silent land: Travels in neuropsychology*. London: Atlantic Books.

Bronfman, Z. Z., Ginsburg, S., and Jablonka, E. (2016). The transition to minimal consciousness through the evolution of associative learning. *Frontiers in Psychology*, 7, article 1954.

Brontë, E. (1847). *Wuthering heights*. London: Thomas Cautley Newby. Full text available at <https://www.gutenberg.org/files/768/768-h/768-h.htm> and <https://books.google.co.uk/books?id=7wXy0iWQhmUC>

Brooks, R. A. (1997). Intelligence without representation. In J. Haugeland (Ed.), *Mind design II: Philosophy, psychology, artificial intelligence* (pp. 395–420). Cambridge, MA: MIT Press. Reprinted (with extra material) from (1991) *Artificial Intelligence*, 47, 139–159.

Brooks, R. A. (2002). *Robot: The future of flesh and machines*. London: Penguin. Also published as *Flesh and machines: How robots will change us*. New York: Pantheon.

Brooks, R. A., Breazeal, C. Marjanović, M., Scassellati, B., and Williamson, M. M. (1998). The cog project: Building a humanoid robot. In C. L. Nehaniv (Ed.), *Computation for metaphors, analogy, and agents. Lecture notes in artificial intelligence* (pp. 52–87). New York: Springer.

Brown, A. (1999). *The Darwin wars: How stupid genes became selfish gods*. London: Simon & Schuster. Also published as *The Darwin wars: The scientific battle for the soul of man*.

Brown, R. (2012). Running on empty: Comments on Prettyman. *Consciousness Online*, 17 February. <https://consciousnessonline.wordpress.com/2012/02/17/empty-thoughts-an-explanatory-problem-for-higher-order-theories-of-consciousness/>

Broyd, S. J., Demanuele, C., Debener, S., Helps, S. K., James, C. J., and Sonuga-Barke, E.J.S. (2009). Default-mode brain dysfunction

• R E F E R E N C E S

in mental disorders: A systematic review. *Neuroscience and Biobehavioral Reviews*, 33, 279–296.

Brugger, P. (2006). From phantom limb to phantom body. In G. Knoblich (Ed.), *Human body perception from the inside out* (pp. 171–209). Oxford: Oxford University Press.

Bugnyar, T., and Heinrich, B. (2005). Ravens, *Corvus corax*, differentiate between knowledgeable and ignorant competitors. *Proceedings of the Royal Society B*, 272, 1641–1646.

Buss, D. M. (1999). *Evolutionary psychology: The new science of the mind*. Boston, MA: Allyn & Bacon.

Buttazzo, G. (2001). Artificial consciousness: Utopia or real possibility? *Computer*, 34(7), 24–30.

Byrne, R. W., and Whiten, A. (Eds) (1988). *Machiavellian intelligence: Social expertise and the evolution of intellect in monkeys, apes and humans*. Oxford: Clarendon Press.

Cairò, O. (2011). External measures of cognition. *Frontiers in Human Neuroscience*, 5, article 108.

Callaway, J. C. (1999). Phytochemistry and neuropharmacology of ayahuasca. In R. Metzner (Ed.), *Sacred vine of spirits: Ayahuasca* (pp. 250–275). New York: Thunder's Mouth Press.

Calvin, W. H. (1996). *The cerebral code: Thinking a thought in the mosaics of the mind*. Cambridge, MA: MIT Press.

Campbell, D. T. (1960). Blind variation and selective retention in creative thought as in other knowledge processes. *Psychological Review*, 67, 380–400.

Campbell, J. (2002). *Reference and consciousness*. Oxford: Oxford University Press.

Campion, J., Latto, R., and Smith, Y. M. (1983). Is blindsight an effect of scattered light, spared cortex, and near-threshold vision? *Behavioral and Brain Sciences*, 6(3), 423–486 (incl. commentaries and authors' response).

Caracciolo, M., and Kukkonen, K. (2014). Cognitive literary study: Second-generation approaches. Special issue, *Style*, 48(3).

Carey, J. M. (2009). *Development and validation of a measure of free will belief and its alternatives*. MA dissertation, University of British Columbia.

Carhart-Harris, R. L., Erritzoe, D., Williams, T., Stone, J. M., Reed, R. J., Colasanti, A., . . . and Nutt, D. J. (2012). Neural correlates of the psychedelic state as determined by fMRI studies with psilocybin. *Proceedings of the National Academy of Sciences of the United States of America*, 109, 2138–2143.

Carhart-Harris, R. L., Bolstridge, M., Rucker, J., Day, C.M.J., Erritzoe, D., Kaelen, M., . . . and Nutt, D. J. (2016a). Psilocybin with

psychological support for treatment-resistant depression: An open-label feasibility study. *Lancet Psychiatry*, 3(7), 619–627.

Carhart-Harris, R. L., Muthukumaraswamy, S., Roseman, L., Kaelen, M., Droog, W., Murphy, K., . . . and Leech, R. (2016b). Neural correlates of the LSD experience revealed by multimodal neuroimaging. *Proceedings of the National Academy of Sciences, USA*, 113(17), 4853–4858.

Carpenter, W. B. (1874). *Principles of mental physiology, with their applications to the training and discipline of the mind and the study of its morbid conditions*. London: Henry S. King & Co.

Carrington, H. (1919). *Modern psychical phenomena: Recent researches and speculations*. London: Kegan Paul, Trench, Trubner & Co.

Carroll, L. (1865). *Alice's adventures in wonderland*. London: Macmillan. Full text available at <https://www.gutenberg.org/files/11/11-h/11-h.htm> and <https://books.google.co.uk/books?id=hWByX5-c5SIC>

Carruthers, P. (2004). Suffering without subjectivity. *Philosophical Studies*, 121, 99–125.

Carruthers, P. (2007). Higher-order theories of consciousness. In M. Velmans and S. Schneider (Eds), *The Blackwell companion to consciousness* (pp. 277–286). Oxford: Blackwell.

Carruthers, P. (2015). Block's overflow argument. *Pacific Philosophical Quarterly*. Online first.

Castaneda, C. (1968). *The teachings of Don Juan: A Yaqui way of knowledge*. University of California Press, Berkeley, CA. Also (1970) London: Penguin.

Castaneda, C. (1971). *A separate reality: Further conversations with Don Juan*. New York: Simon & Schuster.

Castiello, U., Paulignan, Y., and Jeannerod, M. (1991). Temporal dissociation of motor responses and subjective awareness: A study in normal subjects. *Brain*, 114, 2639–2655.

Chagnon, N. A. (1992). *Yanomamö*. 4th ed. Orlando, FL: Harcourt Brace, Jovanovich.

Chalmers, D. J. (1993/2011). A computational foundation for the study of cognition. *Journal of Cognitive Science*, 12, 325–359.

Chalmers, D. J. (1995a). Facing up to the problem of consciousness. *Journal of Consciousness Studies*, 2, 200–219. Reprinted in J. Shear (Ed.). (1997). *Explaining consciousness – The 'hard problem'* (pp. 9–30). Cambridge, MA: MIT Press.

Chalmers, D. J. (1995b). The puzzle of conscious experience. *Scientific American*, December, 62–68.

Chalmers, D. J. (1996). *The conscious mind: In search of a fundamental theory*. Oxford: Oxford University Press.

• R E F E R E N C E S

Chalmers, D.J. (1997). An exchange with David Chalmers. In J. Searle (Ed.), *The mystery of consciousness* (pp. 163–167). New York: New York Review of Books.

Chalmers, D.J. (1999). First-person methods in the science of consciousness. *Consciousness Bulletin*, University of Arizona, June. <http://consc.net/papers/firstperson.html>

Chalmers, D.J. (2000). What is a neural correlate of consciousness? In T. Metzinger (Ed.), *Neural correlates of consciousness* (pp. 17–39). Cambridge, MA: MIT Press.

Chalmers, D.J. (Ed.) (2002). *Philosophy of mind: Classical and contemporary readings*. New York: Oxford University Press.

Chalmers, D.J. (2007). Naturalistic dualism. In M. Velmans and S. Schneider (Eds), *The Blackwell companion to consciousness* (pp. 359–368). Malden, MA: Blackwell.

Chalmers, D.J. (2010). *The character of consciousness*. New York: Oxford University Press.

Chalmers, D.J. (2017). Panpsychism and panprotopsyism. In G. Brüntrup and L. Jaskolla (Eds), *Panpsychism: Contemporary perspectives* (pp. 19–47). New York: Oxford University Press.

Chapman, C.R., and Nakamura, Y. (1999). A passion for the soul: An introduction to pain for consciousness researchers. *Consciousness and Cognition*, 8, 391–422.

Chawla, L.S., Akst, S., Junker, C., Jacobs, B., and Seneff, M.G. (2009). Surges of electroencephalogram activity at the time of death: A case series. *Palliative Medicine*, 12(12), 1095–1100.

Cheesman, J., and Merikle, P.M. (1984). Priming with and without awareness. *Perception and Psychophysics*, 36, 387–395.

Cheesman, J. and Merikle, P.M. (1986). Distinguishing conscious from unconscious perceptual processes. *Canadian Journal of Psychology*, 40, 343–367.

Chella, A., and Manzotti, R. (Eds) (2007). *Artificial consciousness*. Exeter: Imprint Academic.

Cheney, D.L. and Seyfarth, R.M. (1990). *How monkeys see the world: Inside the mind of another species*. Chicago: CV Press.

Chesters, T. (2014). Social cognition: A literary perspective. *Paragraph*, 37(1), 62–78.

Cheyne, J.A., Newby-Clark, I.R., and Rueffer, S.D. (1999). Sleep paralysis and associated hypnagogic and hypnopompic experiences. *Journal of Sleep Research*, 8, 313–317.

Chiang, T. (2005). What's expected of us. *Nature*, 436(7047), 150.

Choi, Y.S., Gray, H.M., and Ambady, N. (2005). The glimpsed world: Unintended communication and unintended perception. In R. R. Hassin, J. S. Uleman, and J. A. Bargh (Eds), *The new unconscious* (pp. 309–333). Oxford: Oxford University Press.

Chopra, D., and Tanzi, R. E. (2012). *Super brain: Unleashing the explosive power of your mind to maximise health, happiness, and spiritual well-being*. New York: Three Rivers.

Chrisley, R. (2009). Artificial intelligence and the study of consciousness. In T. Bayne, A. Cleeremans, and P. Wilken (Eds), *The Oxford companion to consciousness* (pp. 62–66). Oxford: Oxford University Press.

Christensen, J.F., Yoshie, M., Di Costa, S., and Haggard, P. (2016). Emotional valence, sense of agency and responsibility: A study using intentional binding. *Consciousness and Cognition*, 43, 1–10.

Christoff, K., Gordon, A.M., Smallwood, J., Smith, R., and Schooler, J. W. (2009). Experience sampling during fMRI reveals default network and executive system contributions to mind wandering. *Proceedings of the National Academy of Sciences*, 106(21), 8719–8724.

Churchland, P. M. (1981). Eliminative materialism and the propositional attitudes. *Journal of Philosophy*, 78, 67–90.

Churchland, P.M. (1985). Reduction, qualia, and the direct introspection of brain states. *The Journal of Philosophy*, 82(1), 8–28.

Churchland, P.S. (1981). On the alleged backwards referral of experiences and its relevance to the mind – body problem. *Philosophy of Science*, 48, 165–181.

Churchland, P.S. (1988). Reduction and the neurobiological basis of consciousness. In A. J. Marcel and E. Bisiach (Eds), *Consciousness in contemporary science* (pp. 273–304). Oxford: Oxford University Press.

Churchland, P.S. (1996). The Hornswoggle problem. *Journal of Consciousness Studies*, 3(5–6), 402–408. Reprinted in Shear, J. (1997), *Explaining consciousness – The ‘hard problem’* (pp. 37–44). Cambridge, MA: MIT Press.

Churchland, P.S. (1998). Brainshy: Nonneural theories of conscious experience. In S. R. Hameroff, A. W. Kaszniak, and A. C. Scott (Eds), *Toward a science of consciousness: The Second Tucson Discussions and Debates* (pp. 109–124). Cambridge, MA: MIT Press.

Churchland, P.S. (2002). *Brain-wise: Studies in neurophilosophy*. Cambridge, MA: MIT Press.

Cicogna, P., and Bosinelli, M. (2001). Consciousness during dreams. *Consciousness and Cognition*, 10, 26–41.

Clark, A. (1997). *Being there: Putting Brain, body, and world together again*. Cambridge, MA: MIT Press.

• R E F E R E N C E S

- Clark, A.** (2008). *Supersizing the mind: Embodiment, action, and cognitive extension*. Oxford: Oxford University Press.
- Clark, A.** (2013). Whatever next? Predictive brains, situated agents, and the future of cognitive science. *Behavioral and Brain Sciences*, 36(3), 181–253 (incl. commentaries and author's response).
- Clark, A.** (2015). *Surfing uncertainty: Prediction, action, and the embodied mind*. New York: Oxford University Press.
- Clark, A. and Chalmers, D.** (1998). The extended mind. *Analysis*, 58, 7–19. Reprinted in Chalmers, D. (2002), *Philosophy of mind: Classical and contemporary readings* (pp. 643–651). New York: Oxford University Press. Also reprinted in Clark, A. (2008), *Supersizing the mind: Embodiment, action, and cognitive extension* (pp. 220–232). Oxford: Oxford University Press.
- Clarke, C.J.S.** (1995). The nonlocality of mind. *Journal of Consciousness Studies*, 2(3), 231–240. Reprinted in Shear J. (1997), *Explaining consciousness – The 'hard problem'* (pp. 165–175). Cambridge, MA: MIT Press.
- Claxton, G.** (1986a). The light's on but there's nobody home: The psychology of no-self. In G. Claxton (Ed.), *Beyond therapy: The impact of Eastern religions on psychological theory and practice* (pp. 49–70). London: Wisdom.
- Claxton, G.** (Ed.) (1986b). *Beyond therapy: The impact of Eastern religions on psychological theory and practice*. London: Wisdom. Reprinted (1996), Sturminster Newton, Dorset: Prism Press.
- Claxton, G.** (1994). *Noises from the darkroom: The science and mystery of the mind*. London: Aquarian.
- Claxton, G.** (1997). *Hare brain, tortoise mind: Why intelligence increases when you think less*. London: Fourth Estate.
- Cleeremans, A.** (Ed.) (2003). *The unity of consciousness: Binding, integration and dissociation*. New York: Oxford University Press.
- Cleeremans, A.** (2008). Consciousness: The radical plasticity thesis. *Progress in Brain Research* 168, 19–34.
- Clifford, W.** (1874/1886). Body and mind. *Fortnightly Review*, 16, 199–245. Reprinted in L. Stephen and F. Pollock (Eds), *Lectures and essays, by the late William Kingdom Clifford* (pp. 243–273). London: Macmillan. (Page numbers are to the 1886 edition.)
- Clowes, R., Torrance, S., and Chrisley, R.** (Eds) (2007). Machine consciousness. Special issue, *Journal of Consciousness Studies*, 14. Also Exeter: Imprint Academic.
- Cohen, D.** (1987). Behaviourism. In R. L. Gregory (Ed.), *The Oxford companion to the mind* (pp. 71–74). Oxford: Oxford University Press.
- Cohen, M.A., Cavanagh, P., Chun, M. M., and Nakayama, K.** (2012). The attentional requirements of consciousness. *Trends in Cognitive Sciences*, 16(8), 411–417.

- Cohen, M. S., Kosslyn, S. M., Breiter, H. C., DiGirolamo, G. J., Thompson, W. L., Anderson, A. K., Bookheimer, S. Y., Rosen, B. R., and Belliveau, J. W.** (1996). Changes in cortical activity during mental rotation: A mapping study using functional MRI. *Brain*, *119*, 89–100.
- Colzato, L. S., van der Wel, P., Sellaro, R., and Hommel, B.** (2016). A single bout of meditation biases cognitive control but not attentional focusing: Evidence from the global-local task. *Consciousness and Cognition*, *39*, 1–7.
- Conan Doyle, A.** (1887). A study in scarlet. *Beeton's Christmas annual*. London: Ward Lock & Co. Full text available at <https://www.gutenberg.org/files/244/244-h/244-h.htm> and <https://books.google.co.uk/books?id=trM8AAAAYAAJ>
- Conan Doyle, A.** (1891). A scandal in Bohemia. *The Strand magazine*, 25 June. Full text available at <https://www.gutenberg.org/files/1661/1661-h/1661-h.htm>
- Conway, M. A.** (2005). Memory and the self. *Journal of Memory and Language*, *53*, 594–628.
- Cornelissen, F. W., Wade, A. R., Vladusich, T., Dougherty, R. F., and Wandell, B. A.** (2006). No fMRI evidence for brightness and color filling-in in early human visual cortex. *Journal of Neuroscience*, *26*, 3634–3641.
- Costall, A.** (2006). 'Introspectionism' and the mythical origins of scientific philosophy. *Consciousness and Cognition*, *15*, 634–654.
- Cotterill, R. M. J.** (1995). On the unity of conscious experience. *Journal of Consciousness Studies*, *2*(4), 290–312.
- Cotterill, R. M. J.** (1998). *Enchanted looms: Conscious networks in brains and computers*. Cambridge: Cambridge University Press.
- Cotterill, R.** (2003). CyberChild: A simulation test-bed for consciousness studies. *Journal of Consciousness Studies*, *10*, 31–45.
- Cowan, J. D.** (1982). Spontaneous symmetry breaking in large scale nervous activity. *International Journal of Quantum Chemistry*, *22*, 1059–1082.
- Crean, R. D., Crane, N. A., and Mason, B. J.** (2011). An evidence based review of acute and long-term effects of cannabis use on executive cognition functions. *Journal of Addiction Medicine*, *5*(1), 1–8.
- Crick, F.** (1994). *The astonishing hypothesis: The scientific search for the soul*. New York: Scribner's.
- Crick, F., and Koch, C.** (1990). Towards a neurobiological theory of consciousness. *Seminars in the Neurosciences*, *2*, 263–275.
- Crick, F., and Koch, C.** (1998). Consciousness and neuroscience. *Cerebral Cortex*, *8*, 97–107. Also reprinted in B. J. Baars, W. P. Banks, and J. B. Newman (Eds), *Essential sources in the scientific study of consciousness* (pp. 35–53). Cambridge, MA: MIT Press.

• R E F E R E N C E S

- Crick, F., and Koch, C.** (2000). The unconscious homunculus. *Neuropsychanalysis*, 2(1), 3–11. Also reprinted in T. Metzinger (Ed.), *Neural correlates of consciousness: Empirical and conceptual questions* (pp. 103–110). Cambridge, MA: MIT Press.
- Crick, F., and Koch, C.** (2003). A framework for consciousness. *Nature Neuroscience*, 6, 119–126.
- Crick, F., and Mitchison, G.** (1983). The function of dream sleep. *Nature*, 304, 111–114.
- Crook, J.** (1980). *The evolution of human consciousness*. Oxford: Clarendon Press.
- Crook, J.** (1990). Meditation and personal disclosure: The Western Zen retreat. In J. Crook and D. Fontana (Eds), *Space in mind: East-West psychology and contemporary Buddhism* (pp. 156–173). London: Element.
- Crook, J., and Fontana, D.** (Eds) (1990). *Space in mind: East-West psychology and contemporary Buddhism*. London: Element.
- Csikszentmihalyi, M.** (1975). *Beyond boredom and anxiety: Experiencing flow in work and play*. San Francisco: Jossey-Bass.
- Csikszentmihalyi, M.** (1993). *The evolving self: A psychology for the third millennium*. New York: HarperCollins.
- Csikszentmihalyi, M., and Csikszentmihalyi, I. S.** (Eds) (1988). *Optimal experience: Psychological studies of flow in consciousness*. Cambridge: Cambridge University Press.
- Curran, H. V., and Morgan, C.** (2000). Cognitive, dissociative and psychotogenic effects of ketamine in recreational users on the night of drug use and 3 days later. *Addiction*, 95(4), 575–590.
- Cytowic, R. E.** (1993). *The man who tasted shapes*. New York: Putnam.
- Cytowic, R. E., and Eagleman, D. M.** (2009). *Wednesday is indigo blue: Discovering the brain of synesthesia*. Cambridge, MA: MIT Press.
- Damasio, A.** (1994). *Descartes' error: Emotion, reason and the human brain*. New York: Putnams.
- Damasio, A.** (1999). *The feeling of what happens: Body, emotion and the making of consciousness*. London: Heinemann.
- Danckert, J. A., Sharif, N., Haffenden, A. M., Schiff, K. C., and Goodale, M. A.** (2002). A temporal analysis of grasping in the Ebbinghaus illusion: Planning versus online control. *Experimental Brain Research*, 144, 275–280.
- Danforth, A. L., Struble, C. M., Yazar-Klosinski, B., and Grob, C. S.** (2016). MDMA-assisted therapy: A new treatment model for social anxiety in autistic adults. *Progress in Neuro-Psychopharmacology & Biological Psychiatry*, 64, 237–249.

Danielson, N. B., Guo, J. N., and Blumenfeld, H. (2011). The default mode network and altered consciousness in epilepsy. *Behavioral Neurology*, 24(1), 55–65.

Darwin, C. (1839/1909). *The voyage of the Beagle: Journal and remarks, 1832-1835*. New York: Collier Press.

Darwin, C. (1859). *On the origin of species by means of natural selection, or The preservation of favoured races in the struggle for life*. London: Murray.

Darwin, C. (1871). *The descent of man, and selection in relation to sex*. London: John Murray.

Darwin, C. (1872). *The expression of the emotions in man and animals*. London: John Murray. Also (1965), Chicago: University of Chicago Press.

da Vinci, L. (1651). *A treatise on painting [Trattato della pittura]*. Paris: Langlois. Full text available at <http://www.gutenberg.org/ebooks/46915> and <https://books.google.co.uk/books?id=2iVFAAAAYAAJ> (trans. J.F. Rigaud); also <https://archive.org/details/trattatopittura01leon> (original Italian)

Davies, M. (2008). Consciousness and explanation. In L. Weiskrantz and M. Davies (Eds), *Frontiers of consciousness: Chichele lectures* (pp. 1–53). Oxford: Oxford University Press.

Davis, J. H., and Vago, D. R. (2013). Can enlightenment be traced to specific neural correlates, cognition, or behavior? No, and (a qualified) yes. *Frontiers in Psychology*, 4, article 870.

Dawkins, M. S. (2008). The science of animal suffering. *Ethology*, 114(10), 937–945.

Dawkins, R. (1976). *The selfish gene*. Oxford: Oxford University Press. (New edition with additional material, 1989.)

Dawkins, R. (1986). *The blind watchmaker: Why the evidence of evolution reveals a universe without design*. London: Longman.

Dawkins, R. (1989). *The extended phenotype: The long reach of the gene*. Oxford: Oxford University Press.

Deacon, T. (1997). *The symbolic species: The co-evolution of language and the human brain*. London: Penguin.

deCharms, R. C., Maeda, F., Glover, G. H., Ludlow, D., Pauly, J. M., Soneki, D., Gabrieli, J. D. E., and Mackey, S. C. (2005). Control over brain activation and pain learned by using real-time functional MRI. *Proceedings of the National Academy of Sciences, USA*, 102(51), 18626–18631.

De Foe, A., van Doorn, G., and Symmons, M. (2012). Auditory hallucinations predict likelihood of out-of-body experience. *Australian Journal of Parapsychology*, 12(1), 59.

• R E F E R E N C E S

- De Haan, B., Morgan, P.S., and Rorden, C.** (2008). Covert orienting of attention and overt eye movements activate identical brain regions. *Brain Research*, 1204, 102–111.
- Dehaene, S.** (2009). Neuronal global workspace. In T. Bayne, A. Cleere-mans, and P. Wilken (Eds), *The Oxford companion to consciousness* (pp. 466–470). Oxford: Oxford University Press.
- Dehaene, S.** (2014). *Consciousness and the brain: Deciphering how the brain codes our thoughts*. New York: Viking Penguin.
- Dehaene, S., Changeux, J. P., Naccache, L., Sackur, J., and Sergant, C.** (2006). Conscious, preconscious, and subliminal processing: A testable taxonomy. *Trends in Cognitive Sciences*, 10(5), 204–211.
- Dehaene, S., and Naccache, L.** (2001). Towards a cognitive neuroscience of consciousness: Basic evidence and a workspace framework. *Cognition*, 79, 1–37.
- Dehaene, S., Naccache, L., Le Clec'H, G., Koechlin, E., Mueller, M., Dehaene-Lambertz, G., van de Moortele, P.-F., and Le Bihan, D.** (1998). Imaging unconscious semantic priming. *Nature*, 395, 597–600.
- Deikman, A.J.** (1966). Deautomatization and the mystic experience. *Psychiatry*, 29, 324–338.
- Deikman, A.J.** (2000). A functional approach to mysticism. *Journal of Consciousness Studies*, 7(11–12), 75–91.
- de Mille, R.** (1976). *Castaneda's journey: The power and the allegory*. Santa Barbara, CA: Capra Press.
- Denis, D., and Poerio, G.L.** (2017). Terror and bliss? Commonalities and distinctions between sleep paralysis, lucid dreaming, and their associations with waking life experiences. *Journal of Sleep Research*, 26(1), 38–47.
- Dennett, D.C.** (1976). Are dreams experiences? *Philosophical Review*, 73, 151–171. Reprinted in D.C. Dennett (1978). *Brainstorms: Philosophical essays on mind and psychology* (pp. 129–148). New York: Penguin.
- Dennett, D.** (1984/2015). *Elbow room: The varieties of free will worth wanting*. Cambridge, MA: MIT Press.
- Dennett, D.C.** (1987). *The intentional stance*. Cambridge, MA: MIT Press.
- Dennett, D.C.** (1988). Quining qualia. In A. J. Marcel and E. Bisiach (Eds), *Consciousness in contemporary science* (pp. 42–77). Oxford: Oxford University Press. Reprinted in D. Chalmers (2002), *Philosophy of mind: Classical and contemporary readings* (pp. 226–246). New York: Oxford University Press.
- Dennett, D.C.** (1991). *Consciousness explained*. Boston, MA: Little, Brown and Co.
- Dennett, D.C.** (1995a). The path not taken. *Behavioral and Brain Sciences*, 18, 252–253. Commentary on N. Block, 'On a confusion about a function of consciousness', *BBS*, 18, 227–247.

Dennett, D. C. (1995b). *Darwin's dangerous idea: Evolution and the meaning of life*. London: Penguin.

Dennett, D. C. (1995c). The unimagined preposterousness of zombies. *Journal of Consciousness Studies*, 2(4), 322–326. Commentary on T. Moody's 'Conversations with zombies'.

Dennett, D. C. (1995d). Cog: Steps towards consciousness in robots. In T. Metzinger (Ed.), *Conscious experience* (pp. 471–487). Thorverton: Imprint Academic.

Dennett, D. C. (1996a). Facing backwards on the problem of consciousness. *Journal of Consciousness Studies*, 3(1), 4–6.

Dennett, D. C. (1996b). *Kinds of minds: Towards an understanding of consciousness*. London: Weidenfeld & Nicolson.

Dennett, D. C. (1997). An exchange with Daniel Dennett. In J. Searle (Ed.), *The mystery of consciousness* (pp. 115–119). New York: New York Review of Books.

Dennett, D. C. (1998a). The myth of double transduction. In R. Hameroff, A. W. Kaszniak, and A. C. Scott (Eds), *Toward a science of consciousness: The second Tucson discussions and debates* (pp. 97–107). Cambridge, MA: MIT Press. [5, p. 22]

Dennett, D. C. (1998b). *Brainchildren: Essays on designing minds*. Cambridge, MA: MIT Press.

Dennett, D. C. (2001a). Are we explaining consciousness yet? *Cognition*, 79(1–2), 221–237.

Dennett, D. C. (2001b). *The fantasy of first person science*. Debate with D. Chalmers, Northwestern University, Evanston, IL, Feb 2001. <http://ase.tufts.edu/cogstud/papers/chalmersdeb3dff.htm>

Dennett, D. C. (2003). *Freedom evolves*. New York: Penguin.

Dennett, D. C. (2005). *Sweet dreams: Philosophical obstacles to a science of consciousness*. Cambridge, MA: MIT Press.

Dennett, D. C. (2007). Heterophenomenology reconsidered. *Phenomenology and Cognitive Science*, 6, 247–270.

Dennett, D. C. (2011). Shall we tango? No, but thanks for asking. *Journal of Consciousness Studies*, 18(5–6), 23–34.

Dennett, D. C. (2014). Reflections on 'free will'. Naturalism.org, 24 January 2014. www.naturalism.org/resources/book-reviews/reflections-on-free-will

Dennett, D. C. (2016). Illusionism as the obvious default theory of consciousness. *Journal of Consciousness Studies*, 23(11–12), 65–72.

Dennett, D. C. (2017). *From bacteria to Bach and back: The evolution of minds*. London: Allen Lane.

• R E F E R E N C E S

- Dennett, D.C., and Kinsbourne, M.** (1992). Time and the observer: The where and when of consciousness in the brain. *Behavioral and Brain Sciences*, 15, 183–247 (incl. commentaries and authors' response).
- Depraz, N.** (1999). The phenomenological reduction as *praxis*. *Journal of Consciousness Studies*, 6(2–3), 95–110. Reprinted in F.J. Varela and J. Shear (Eds) (1999). *The view from within* (pp. 95–110). Thorverton, Devon: Imprint Academic.
- De Ridder, D., Van Laere, K., Dupont, P., Menovsky, T., and Van de Heyning, P.** (2007). Visualizing out-of-body experience in the brain. *New England Journal of Medicine*, 357, 1829–1833.
- Descartes, R.** (1637/1649). *A discourse of a method for the well-guiding of reason, and the discovery of truth in the sciences* [*Discours de la méthode pour bien conduire sa raison, et chercher la vérité dans les sciences*]. London: Thomas Newcombe, for John Holden.
- Descartes, R.** (1641/2008). *Meditations on first philosophy* [*Meditationes de prima philosophia*]. Trans. J. Veitch. New York: Cosimo.
- de'Sperati, C., and Santandrea, E.** (2005). Smooth pursuit-like eye movements during mental extrapolation of motion: The facilitatory effect of drowsiness. *Cognitive Brain Research*, 25, 328–338.
- Devereux, P.** (1997). *The long trip: A prehistory of psychedelia*. London: Penguin.
- De Weerd, P., Gattas, R., Desimone, R., and Ungerleider, L.G.** (1995). Responses of cells in monkey visual cortex during perceptual filling-in of an artificial scotoma. *Nature*, 377, 731–734.
- Di, H., Boly, M., Weng, X., Ledoux, D., and Laureys, S.** (2008). Neuroimaging activation studies in the vegetative state: Predictors of recovery? *Clinical Medicine*, 8(5), 502–507.
- Dickinson, E.** (1999). *The poems of Emily Dickinson: Reading edition*. Ed. R.W. Franklin. Cambridge, MA: The Belknap Press of Harvard University Press.
- Diderot, D.** (1769). Conversation between d'Alembert and Diderot [Entretien entre d'Alembert et Diderot]. In *Le Rêve de d'Alembert* [D'Alembert's dream]. Full text available at <https://www.marxists.org/reference/archive/diderot/1769/conversation.htm>; also <https://archive.org/details/entretienentreda00dideuoft> (French original)
- Dietrich, A.** (2007). *Introduction to consciousness*. New York: Palgrave Macmillan.
- Dijksterhuis, A., Aarts, H., and Smith, P.K.** (2005). The power of the subliminal: On subliminal persuasion and other potential applications. In R.R. Hassin, J.S. Uleman, and J.A. Bargh (Eds), *The new unconscious* (pp. 77–106). Oxford: Oxford University Press.
- Dixon, N.F.** (1971). *Subliminal perception: The nature of a controversy*. London: McGraw-Hill.

- Doblin, R.** (1991). Pahnke's 'Good Friday Experiment': A long-term follow-up and methodological critique. *The Journal of Transpersonal Psychology*, 23, 1–28.
- Dobzhansky, T.** (1973). Nothing in biology makes sense except in the light of evolution. *American Biology Teacher*, 35(3), 125–129.
- Domhoefer, S.M., Unema, P.J.A., and Velichkovsky, B.M.** (2002). Blinks, blanks and saccades: How blind we really are for relevant visual events. *Progress in Brain Research*, 140, 119–131.
- Domhoff, G.W.** (1996). *Finding meaning in dreams: A quantitative approach*. New York: Plenum Press.
- Donald, M.** (1991). *Origins of the modern mind: Three stages in the evolution of culture and cognition*. Cambridge, MA: Harvard University Press.
- Donald, M.** (2001). *A mind so rare: The evolution of human consciousness*. New York: W. W. Norton.
- Dong, Y., Mihalas, S., Qiu, F., von der Heydt, R., and Niebur, E.** (2008). Synchrony and the binding problem in macaque visual cortex. *Journal of Vision*, 8(7), 1–16.
- dos Santos, R.G., Osorio, F.L., Crippa, J.A.S., Riba, J., Zuardi, A.W., and Hallak, J.E.C.** (2016). Antidepressive, anxiolytic, and antiaddictive effects of ayahuasca, psilocybin and lysergic acid diethylamide (LSD): A systematic review of clinical trials published in the last 25 years. *Therapeutic Advances in Psychopharmacology*, 6(3), 193–213.
- Dostoyevsky, F.** (1864). *Notes from the underground* [Записки из подполья]. *Epoch*, March. Full text available at <https://www.gutenberg.org/files/600/600-h/600-h.htm>; also <http://rvb.ru/dostoevski/01text/vol4/24.htm> (original Russian)
- Dresler, M., Wehrle, R., Spoormaker, V.I., Koch, S.P., Holsboer, F., Steiger, A., . . . and Czisch, M.** (2012). Neural correlates of dream lucidity obtained from contrasting lucid versus non-lucid dream sleep: A combined EEG/fMRI case study. *Sleep*, 35(7), 1017–1020.
- Dumas, A.** (1846). *The Count of Monte Cristo* [Le comte de Monte Cristo]. *Journal des débats*, 28 August 1844 - 15 January 1846; *L'Écho des Feuilles*, 1846. Full text available at <http://www.gutenberg.org/ebooks/1184> and <https://books.google.co.uk/books?id=RyEEAAAQAAJ>; also <http://www.gutenberg.org/ebooks/17989> (original French, vol. 1)
- Dunbar, E.** (1905). The light thrown on psychological processes by the action of drugs. *Proceedings of the Society for Psychical Research*, 19, 62–77.
- Dunbar, R.** (1996). *Grooming, gossip and the evolution of language*. London: Faber & Faber.
- Dunning, A., and Woodrow, P.** (2010). Machine imagination: Closed eye hallucination and the ganzfeld effect. *13th Generative Art Conference GA2010*, 35–45.

• R E F E R E N C E S

- Eagleman, D.M., and Holcombe, A.O.** (2002). Causality and the perception of time. *Trends in Cognitive Sciences*, 6, 323–325.
- Eagleman, D.M., and Sejnowski, T.J.** (2000). Motion integration and postdiction in visual awareness. *Science*, 287(5460), 2036–2038.
- Earleywine, M.** (2002). *Understanding marijuana: A new look at the scientific evidence*. New York: Oxford University Press.
- Eccles, J.C.** (1994). *How the self controls its brain*. Berlin: Springer.
- Edelman, D.B., and Seth, A.K.** (2009). Animal consciousness: A synthetic approach. *Trends in Neurosciences*, 32(9), 476–484.
- Edelman, D.B., Baars, B.J., and Seth, A.K.** (2005). Identifying hallmarks of consciousness in non-mammalian species. *Consciousness and Cognition*, 14(1), 169–187.
- Edelman, G.M.** (1989). *Neural Darwinism: The theory of neuronal group selection*. Oxford: Oxford University Press.
- Edelman, G.M.** (2003). Naturalizing consciousness: A theoretical framework. *Proceedings of the National Academy of Sciences, USA*, 100(9), 5520–5524.
- Edelman, G.M., and Tononi, G.** (2000a). *Consciousness: How matter becomes imagination*. London: Penguin. Also published as (2000a) *A universe of consciousness: How matter becomes imagination*. New York: Basic Books.
- Edelman, G.M., and Tononi, G.** (2000b). Reentry and the dynamic core: Neural correlates of conscious experience. In T. Metzinger (Ed.), *Neural correlates of consciousness* (pp. 139–151). Cambridge, MA: MIT Press.
- Ehrsson, H.** (2007). The experimental induction of out-of-body experiences. *Science*, 317(5841), 1048.
- Einstein, A.** (1930). ‘What I believe’. *Forum and Century* (1930–1940), October, LXXXIV(4), 192.
- Einstein, A.** (1949/2006). *The world as I see it [Mein Weltbild]*. Trans. A. Harris. New York: Kensington.
- Eklund, A., Nichols, T.E., and Knutsson, H.** (2016). Cluster failure: Why fMRI inferences for spatial extent have inflated false-positive rates. *Proceedings of the National Academy of Sciences, USA*, 113(28), 7900–7905.
- Elwood, R.W., Barr, S., and Patterson, L.** (2009). Pain and stress in crustaceans? *Applied Animal Behaviour Science*, 118(3–4), 128–136.
- Empson, J.** (2001). *Sleep and dreaming*. 3rd ed. New York: Palgrave Macmillan.
- Engel, A. K.** (2003). Temporal binding and the neural correlates of consciousness. In A. Cleeremans (Ed.), *The unity of consciousness: Binding, integration and dissociation* (pp. 132–152). New York: Oxford University Press.

- Engel, A.K., Fries, P., König, P., Brecht, M., and Singer, W.** (1999). Temporal binding, binocular rivalry, and consciousness. *Consciousness and Cognition*, 8, 128–151.
- Engler, J.** (1986). Therapeutic aims in psychotherapy and meditation. In K. Wilber, J. Engler, and D. Brown (Eds), *Transformations of consciousness: Conventional and contemplative perspectives on development* (pp. 17–51). Boston, MA: Shambhala.
- Engler, J.** (2003). Being somebody and being nobody: A re-examination of the understanding of self in psychoanalysis and Buddhism. In J.D. Safran (Ed.), *Psychoanalysis and Buddhism: An unfolding dialogue* (pp. 35–79). Boston: Wisdom.
- Erasmus, D.** (1524/1999). *Discourse on free will [De libero arbitrio diatribesive collatio]*. Trans. E.F. Winter. New York: Continuum.
- Eriksen, C.W., and St James, J.D.** (1986). Visual attention within and around the field of focal attention: A zoom lens model. *Perception & Psychophysics*, 40(4), 225–240.
- Erlacher, D., Schädlich, M., Stumbrys, T., and Schredl, M.** (2014). Time for actions in lucid dreams: Effects of task modality, length, and complexity. *Frontiers in Psychology*, 4, article 1013.
- Erlacher, D., and Schredl, M.** (2004). Time required for motor activity in lucid dreams. *Perceptual and Motor Skills*, 99, 1239–1242.
- Erlacher, D., and Schredl, M.** (2008). Do REM (lucid) dreamed and executed actions share the same neural substrate? *International Journal of Dream Research*, 1, 7–14.
- Erlacher, D., Stumbrys, T., and Schredl, M.** (2012). Frequency of lucid dreams and lucid dream practice in German athletes. *Imagination, Cognition and Personality*, 31(3), 237–246.
- Evans, K.K., Horowitz, T.S., Howe, P., Pedersini, R., Reijnen, E., Pinto, Y., Kuzmova, Y., and Wolf, J.M.** (2011). Visual attention. *Wiley Interdisciplinary Reviews: Cognitive Science*, 2, 503–514.
- Faraday, M.** (1853). Experimental investigations of table-moving. *The Athenaeum*, 1340, 801–803.
- Farrell, B.** (1996). Review of *The body and the self*, ed. J.L. Bermudez, A. Marcel, and N. Eilan. *Journal of Consciousness Studies*, 3, 517–519.
- Farrow, J.T., and Hebert, J.R.** (1982). Breath suspension during the transcendental meditation technique. *Psychosomatic medicine*, 44, 133–153.
- Farthing, G.W.** (1992). *The psychology of consciousness*. Englewood Cliffs, NJ: Prentice Hall.
- Faymonville, M.E., Laureys, S., Degueldre, C., Del Flore, G., Luxen, A., Franck, G., Lamy, M., and Maquet, P.** (2000). Neural

• R E F E R E N C E S

mechanisms of antinociceptive effects of hypnosis. *Anesthesiology*, 92, 1257–1267.

Feinberg, T. E. (2001). Why the mind is not a radically emergent feature of the brain. *Journal of Consciousness Studies*, 8, 123–145. Also reprinted in A. Freeman (Ed.), *The emergence of consciousness* (pp. 123–145). Charlottesville, VA: Imprint Academic.

Feinberg, T. E. (2009). *From axons to identity: Neurological explorations of the nature of self*. New York: Norton.

Feinberg, T. E., and Mallatt, J. M. (2016). *The ancient origins of consciousness: How the brain created consciousness*. Cambridge, MA: MIT Press.

Fell, J., Axmacher, N., and Haupt, S. (2010). From alpha to gamma: Electrophysiological correlates of meditation-related states of consciousness. *Medical Hypotheses*, 75, 218–224.

Feltz, A., and Cova, F. (2014). Moral responsibility and free will: A meta-analysis. *Consciousness and Cognition*, 30, 234–246.

Fénelon, G., Mahieux, F., Huon, R., and Ziegler, M. (2000). Hallucinations in Parkinson's disease: Prevalence, phenomenology and risk factors. *Brain*, 123(4), 733–745.

Fenwick, P. (1987). Meditation and the EEG. In M. West (Ed.), *The psychology of meditation* (pp. 104–117). Oxford: Clarendon Press.

Feynman, R., and Leighton, R. (1985). *'Surely you're joking Mr. Feynman!'* *Adventures of a curious character*. New York: W.W. Norton.

ffytche, D. H. (2000). Imaging conscious vision. In T. Metzinger (Ed.), *Neural correlates of consciousness* (pp. 221–230). Cambridge, MA: MIT Press.

ffytche, D. H., and Howard, R. J. (1999). The perceptual consequences of visual loss: 'Positive' pathologies of vision. *Brain*, 122, 1247–1260.

ffytche, D. H., Howard, R. J., Brammer, M. J., David, A., Woodruff, P., and Williams, S. (1998). The anatomy of conscious vision: An fMRI study of visual hallucinations. *Nature Neuroscience*, 1, 738–742.

Fichte, J. G. (1794–1795/1982). *The science of knowledge [Grundlage der gesamten Wissenschaftslehre]*. Trans. P. Heath and J. Lachs. Cambridge: Cambridge University Press.

Filevich, E., Vanneste, P., Brass, M., Fias, W., Haggard, P., and Kühn, S. (2013). Brain correlates of subjective freedom of choice. *Consciousness and Cognition*, 22, 1271–1284.

Finn, E. S., Shen, X., Scheinost, D., Rosenberg, M. D., Huang, J., Chun, M. M., Papademetris, X., Constable, R. T. (2015). Functional connectome fingerprinting: Identifying individuals using patterns of brain connectivity. *Nature Neuroscience*, 18(11), 1665–1671.

- Fisher, M.P.A.** (2015, 29 August). Quantum cognition: The possibility of processing with nuclear spins in the brain. *Annals of Physics*, 362, 593–602. arXiv:1508.05929.
- Fitzgerald, F. S.** (1934). *Tender is the night*. New York: Charles Scribner's Sons. Full text available at <http://gutenberg.net.au/ebooks03/0301261h.html>
- Flanagan, O.** (1992). *Consciousness reconsidered*. Cambridge, MA: MIT Press.
- Flanagan, O.** (2000). *Dreaming souls: Sleep, dreams, and the evolution of the conscious mind*. New York: Oxford University Press.
- Flanagan, O. and Polger, T.** (1995). Zombies and the function of consciousness. *Journal of Consciousness Studies*, 2(4), 313–321.
- Flaubert, G.** (1856). *Madame Bovary*. Paris: Michel Lévy Frères. Full text available at <https://www.gutenberg.org/files/2413/2413-h/2413-h.htm> (trans. E. Marx-Aveling); also http://flaubert.univ-rouen.fr/bovary/bovary_6/doc0/roman.html (original French)
- Flaubert, G.** (1869). *Sentimental education; Or, the history of a young man (L'Éducation sentimentale, histoire d'un jeune homme)*. Paris: Michel Lévy Frères. Full text available at <http://www.gutenberg.org/ebooks/34828> (vol. 1) and <http://www.gutenberg.org/ebooks/27537> (vol. 2); also <http://gallica.bnf.fr/ark:/12148/bpt6k691688> (original French)
- Fletcher, J.A., and Doebeli, M.** (2009). A simple and general explanation for the evolution of altruism. *Proceedings of the Royal Society B*, 276, 13–19.
- Fletcher, P.** (2002). *Seeing with sound: A journey into sight*. Paper presented at Toward a Science of Consciousness, Tucson, AZ, April 8–12, 2002. Conference Research Abstracts (provided by *Journal of Consciousness Studies*), Abstract No. 188.
- Flohr, H.** (2000). NMDA receptor-mediated computational processes and phenomenal consciousness. In T. Metzinger (Ed.), *Neural correlates of consciousness* (pp. 245–258). Cambridge, MA: MIT Press.
- Foglia, L., and O'Regan, J.K.** (2015). A new imagery debate: Enactive and sensorimotor accounts. *Review of Philosophy and Psychology*, 7(1), 181–196.
- Forman, R.K.C.** (Ed.) (1990). *The problem of pure consciousness: Mysticism and philosophy*. New York: Oxford University Press.
- Forman, R.K.C.** (1999). *Mysticism, mind, consciousness*. Albany, NY: State University of New York Press.
- Fowles, J.** (1966/2010). *The magus*. London: Jonathan Cape. / London: Vintage.
- Fowles, J.** (1969/2004). *The French lieutenant's woman*. London: Jonathan Cape. / London: Vintage.

- Fox, K.C.R., Nijeboer, S., Dixon, M.L., Floman, J.L., Ellamil, M., Rumak, S.P., Sedlmeier, P., and Christoff, K.** (2014). Is meditation associated with altered brain structure? A systematic review and meta-analysis of morphometric neuroimaging in meditation practitioners. *Neuroscience and Biobehavioral Reviews*, 43, 48–73.
- Fox, K.C.R., Kang, Y., Lifshitz, M., and Christoff, K.** (2016). Increasing cognitive-emotional flexibility with meditation and hypnosis: The cognitive neuroscience of de-automatization. In A. Raz and M. Lifshitz (Eds), *Hypnosis and meditation* (pp. 191–219). New York: Oxford University Press.
- Fox, M.D., Corbetta, M., Snyder, A.Z., Vincent, J.L., and Raichle, M.E.** (2006). Spontaneous neuronal activity distinguishes human dorsal and ventral attention systems. *Proceedings of the National Academy of Sciences*, 103, 10046–10051.
- Fox, O.** (1962). *Astral projection: A record of out-of-the-body experiences*. New York: University Books.
- Foulkes, D.** (1993). Children’s dreaming. In C. Cavallero and D. Foulkes (Eds), *Dreaming as cognition* (pp. 114–132). New York: Harvester Wheatsheaf.
- Francescotti, R.** (2016). Supervenience and mind. *The Internet Encyclopedia of Philosophy*, ISSN 2161–0002, <http://www.iep.utm.edu/supermin/>, 21 May 2016.
- Francis, B.** (2014). *The robot rendezvous problem*. Bode Lecture at 53rd IEEE Conference on Decision and Control. www.ieeecss-oll.org/node/88
- Frank, M.G., Waldrop, R.H., Dumoulin, M., Aton, S., and Boal, J.G.** (2012). A preliminary analysis of sleep-like states in the cuttlefish *Sepia officinalis*. *PLOS ONE*, 7(6), e38125.
- Frankish, C.** (2016a). Illusionism. Special issue, *Journal of Consciousness Studies*, 23(11–12).
- Frankish, C.** (2016b). Illusionism as a theory of consciousness. *Journal of Consciousness Studies*, 23(11–12), 11–39.
- Frankish, K., and Evans, J. St. B. T.** (2009). The duality of mind: An historical perspective. In J. St. B. T. Evans and K. Frankish (Eds), *In two minds: Dual processes and beyond* (pp. 1–29). New York: Oxford University Press.
- Franklin, S.** (2003). A conscious artifact? *Journal of Consciousness Studies*, 10(4–5), 47–66.
- Franklin, S., and Patterson, F.G., Jr.** (2006). The LIDA architecture: Adding new modes of learning to an intelligent, autonomous, software agent. *Integrated Design and Process Technology*, IDPT-2006 Proceedings.
- Franz, V.H., Gegenfurtner, K.R., Bühlhoff, H.H., and Fahle, M.** (2000). Grasping visual illusions: No evidence for a dissociation between perception and action. *Psychological Science*, 11, 20–25.

Frecska, E., Bokor, P., and Winkelmann, M. (2016). The therapeutic potentials of ayahuasca: Possible effects against various diseases of civilization. *Frontiers in Pharmacology*, 7, article 35.

Freeman, D., Antley, A., Ehlers, A., Dunn, G., Thompson, C., Vorontsova, N., Garety, P., Kuipers, E., Glucksman, E., and Slater, M. (2014). The use of immersive virtual reality (VR) to predict the occurrence 6 months later of paranoid thinking and posttraumatic stress symptoms assessed by self-report and interviewer methods: A study of individuals who have been physically assaulted. *Psychological Assessment*, 26(3), 841–847.

Frege, G. (1918/1967). The thought: A logical inquiry. In P.F. Strawson (ed.), *Philosophical logic* (pp. 17–38). Oxford: Oxford University Press.

Freud, S. (1900/1999). *The interpretation of dreams [Die Traumdeutung]*. Trans. J. Crick. Oxford: Oxford University Press.

Freud, S. (1915). The unconscious. In *General psychological theory: Papers on metapsychology* (pp. 116–150). New York: Collier. www.sas.upenn.edu/~cavitch/pdf-library/Freud_Unconscious.pdf

Freud, S. (1923/1927). *The ego and the id [Das Ich und das Es]*. Trans. J. Riviere. London: Hogarth Press; Institute of Psychoanalysis.

Frijda, N.H. (2007). *The laws of emotion*. Mahwah, NJ: Lawrence Erlbaum.

Frith, C.D. (2007). *Making up the mind: How the brain creates our mental world*. Oxford: Blackwell.

Frith, C.D. (2015). *The cognitive neuropsychology of schizophrenia*. Classic ed. Hove, East Sussex: Psychology Press.

Frith, C.D., Friston, K. Liddle, P.F., and Frakowiak, R.S.J. (1991). Willed action and the prefrontal cortex in man: A study with PET. *Proceedings of the Royal Society B*, 244, 241–246.

Frith, C.D., and Metzinger, T. (2016). What's the use of consciousness? How the stab of conscience made us really conscious. In A. K. Engel, K.J. Friston, and D. Kragic (Eds), *The pragmatic turn: Toward action-oriented views in cognitive science* (pp. 193–214). Cambridge, MA: MIT Press.

Frith, C.D., and Paulesu, E. (1997). The physiological basis of synaesthesia. In S. Baron-Cohen and J.E. Harrison (Eds), *Synaesthesia: Classics and contemporary readings* (pp. 123–147). Oxford: Blackwell.

Froese, T., Gould, C., and Barrett, A. (2011). Re-viewing from within: A commentary on first- and second-person methods in the science of consciousness. *Constructivist Foundations*, 6(2), 254–269.

Froese, T., Iizuka, H., and Ikegami, T. (2014). Embodied social interaction constitutes social cognition in pairs of humans: A minimalist virtual reality experiment. *Scientific Reports*, 4, 3672.

• R E F E R E N C E S

- Fuchs, I., Ansoorge, U., Huber-Huber, C., Höflich, A., and Lanzenberger, R.** (2015). S-ketamine influences strategic allocation of attention but not exogenous capture of attention. *Consciousness and Cognition*, 35, 282–294.
- Fulkerson, M.** (2014). Rethinking the senses and their interactions: The case for sensory pluralism. *Frontiers in Psychology*, 145, article 1426.
- Fuller, S., and Carrasco, M.** (2006). Exogenous attention and color perception: Performance and appearance of saturation and hue. *Vision Research*, 46(23), 4032–4047.
- Gabbard, G.O., and Twemlow, S.W.** (1984). *With the eyes of the mind: An empirical analysis of out-of-body states*. New York: Praeger.
- Gackenbach, J., and Bosveld, J.** (1989). *Control your dreams: How lucid dreaming can help you uncover your hidden fears & explore the frontiers of human consciousness*. New York: Harper & Row.
- Gackenbach, J., and LaBerge, S.** (Eds) (1988). *Conscious mind, sleeping brain: Perspectives on lucid dreaming*. New York: Plenum.
- Gaillard, R., Dehaene, S., Adam, C., Clémenceau, S., Hasboun, D., Baulac, M., Cohen, L., and Naccache, L.** (2009). Converging intracranial markers of conscious access. *PLOS Biology*, 7(3), e1000061.
- Galak, J., Leboeuf, R.A., Nelson, L.D., and Simmons, J.P.** (2012). Correcting the past: Failures to replicate psi. *Journal of Personality and Social Psychology*, 103(6), 933–948.
- Gallagher, S.** (2005). *How the body shapes the mind*. Oxford: Oxford University Press.
- Gallagher, S.** (2007). Phenomenological approaches to consciousness. In M. Velmans and S. Schneider (Eds), *The Blackwell companion to consciousness* (pp. 686–696). Oxford: Blackwell.
- Gallagher, S.** (2008). *Brainstorming: Views and interviews on the mind*. Exeter, United Kingdom: Imprint Academic.
- Gallagher, S.** (2012). *Phenomenology*. Basingstoke: Palgrave Macmillan.
- Gallagher, S., and Zahavi, D.** (2012). *The phenomenological mind*. 2nd ed. London: Routledge.
- Gallimore, A.R.** (2015). Restructuring consciousness – the psychedelic state in light of integrated information theory. *Frontiers in Human Neuroscience*, 9, article 346.
- Gallup, G.G.** (1970). Chimpanzees: Self-recognition. *Science*, 167, 86–87.
- Gallup, G.G.** (1998). Can animals empathize? Yes. *Scientific American*, 9(4), 67–76.
- Gallup, G.H., and Newport, F.** (1991). Belief in paranormal phenomena among adult Americans. *Skeptical Inquirer*, 15, 137–146.

- Galpin, A., Underwood, G., and Crundall, D.** (2009). Change blindness in driving scenes. *Transportation Research Part F*, *12*, 179–185.
- Galton, F.** (1883). *Inquiries into human faculty and its development*. London: Macmillan.
- Garrison, J.R., Bond, R., Gibbard, E., Johnson, M.K., and Simons, J. S.** (2017). Monitoring what is real: The effects of modality and action on accuracy and type of reality monitoring error. *Cortex*, *87*, 108–117.
- Garrison, K.A., Santoyo, J.F., Davis, J.H., Thornhill, T.A., Kerr, C.E., and Brewer, J. A.** (2013). Effortless awareness: Using real time neurofeedback to investigate correlates of posterior cingulate cortex activity in meditators' self-report. *Frontiers in Human Neuroscience*, *7*, article 440.
- Gaspar, J.G., Street, W.N., Windsor, M.B., Carbonari, R., Kaczmariski, H., Kramer, A.F., and Mathewson, K.E.** (2014). Providing views of the driving scene to drivers' conversation partners mitigates cell-phone-related distraction. *Psychological Science*, *25*(12), 2136–2146.
- Gasser, P., Kirchner, K., and Passie, T.** (2014). LSD-assisted psychotherapy for anxiety associated with a life-threatening disease: A qualitative study of acute and sustained subjective effects. *Journal of Psychopharmacology*, *29*(1), 57–68.
- Gauld, A.** (1968). *The founders of psychical research*. London: Routledge and Kegan Paul.
- Gazzaniga, M.S.** (1992). *Nature's mind: The biological roots of thinking, emotions, sexuality, language, and intelligence*. London: Basic Books.
- Geldard, F.A., and Sherrick, C. E.** (1972). The cutaneous 'rabbit': A perceptual illusion. *Science*, *178*, 178–179.
- Gennaro, R.J.** (Ed.) (2004). *Higher-order theories of consciousness: An anthology*. Amsterdam: John Benjamins.
- Gennaro, R.J.** (2017). *Consciousness*. New York: Routledge.
- Geraerts, E., Bernstein, D.M., Merckelbach, H., Linders, C., Raymaekers, L., and Loftus, E. F.** (2008). Lasting false beliefs and their behavioral consequences. *Psychological Science*, *19*, 749–753.
- Gergen, K.J.** (2011). The social construction of self. In S. Gallagher (Ed.), *The Oxford handbook of the self* (pp. 633–653). New York: Oxford University Press.
- Gibson, J. J.** (1979). *The ecological approach to visual perception*. New York: Houghton Mifflin.
- Glimcher, P. W., and Fehr, E.** (Eds.) (2013). *Neuroeconomics: Decision making and the brain*. 2nd ed. Amsterdam: Academic Press.
- Glover, S.** (2002). Visual illusions affect planning but not control. *Trends in Cognitive Sciences*, *6*, 288–292.

• R E F E R E N C E S

Godfrey-Smith, P. (2016). *Other minds: The octopus, the sea, and the deep origins of consciousness*. New York: Farrar, Straus and Giroux. Also published (2017) as *Other minds: The octopus, the sea, and the evolution of intelligent life*. London: William Collins.

Godfrey-Smith, P. (2017). The mind of an octopus. *Scientific American*, 1 January. www.scientificamerican.com/article/the-mind-of-an-octopus/

Goldberg, I. I., Harel, M., and Malach, R. (2006). When the brain loses its self: Prefrontal inactivation during sensorimotor processing. *Neuron*, 50(2), 329–339.

Goodale, M. A. (2007). Duplex vision: Separate cortical pathways for conscious perception and the control of action. In M. Velmans and S. Schneider (Eds), *The Blackwell companion to consciousness* (pp. 616–627). Oxford: Blackwell.

Goodale, M. A. (2014). How (and why) the visual control of action differs from visual perception. *Proceedings of the Royal Society of London B: Biological Sciences*, 281(1785), 20140337.

Goodale, M. A., and Milner, D. (2013). *Sight unseen: An exploration of conscious and unconscious vision*. 2nd ed. Oxford: Oxford University Press.

Goodale, M. A., Pelisson, D., and Prablanc, C. (1986). Large adjustments in visually guided reaching do not depend on vision of the hand or perception of target displacement. *Nature*, 320, 748–750.

Gopnik, A. (2016). How animals think: A new look at what humans can learn from nonhuman minds. *The Atlantic*, May. www.theatlantic.com/magazine/archive/2016/05/how-animals-think/476364/

Goswami, A. (2008). *Creative evolution: A physicist's resolution between Darwinism and intelligent design*. Wheaton, IL: Quest Books.

Gould, S. J., and Lewontin, R. C. (1979). The spandrels of San Marco and the Panglossian paradigm: A critique of the adaptationist program. *Proceedings of the Royal Society B*, 205, 581–598.

Gray, J. (2004). *Consciousness: Creeping up on the hard problem*. Oxford: Oxford University Press.

Gray, J. (2008). The atheist delusion. *The Guardian*, 15 March 2008. www.theguardian.com/books/2008/mar/15/society

Graziano, M. (2013). Consciousness and the unashamed rationalist. *Huffington Post*, 30 August 2013. www.huffingtonpost.com/michael-graziano/consciousness-and-the-una_b_3844493.html

Graziano, M. (2016). Consciousness engineered. *Journal of Consciousness Studies*, 23(11–12), 98–115.

Graziano, M. S., and Kastner, S. (2011). Human consciousness and its relationship to social neuroscience: A novel hypothesis. *Cognitive Neuroscience*, 2(2), 98–113.

- Green, C.E.** (1968a). *Lucid dreams*. London: Hamish Hamilton.
- Green, C.E.** (1968b). *Out-of-the-body experiences*. London: Hamish Hamilton.
- Green, C.E., and McCreery, C.** (1975). *Apparitions*. London: Hamish Hamilton.
- Greene, J., and Cohen, J.** (2004). For the law, neuroscience changes nothing and everything. *Philosophical Transactions of the Royal Society B*, 359(1451), 1775–1785. Also reprinted in M. Tonry (Ed.) (2011), *Why punish? How much? A reader on punishment* (pp. 293-314). New York: Oxford University Press.
- Greenfield, S.** (2000). *Brain story: Why do we think and feel as we do?* London: BBC.
- Gregory, R.L.** (1966/1997). *Eye and brain: The psychology of seeing*. 5th ed. London: Weidenfeld & Nicolson.
- Gregory, R.L.** (1986). *Odd perceptions*. London: Routledge.
- Gregory, R.L.** (1990). Personal communication. (This quip has subsequently been attributed to many people but Richard was a great punner and Sue recalls him saying this to her in the lift in the medical school at Bristol University in 1990.)
- Gregory, R.L.** (2004). *The Oxford companion to the mind*. 2nd ed. Oxford: Oxford University Press.
- Greyson, B.** (2003). Incidence and correlates of near death experiences in a cardiac care unit. *General Hospital Psychiatry*, 25(4), 269–276.
- Griffin, D.R., and Speck, G. B.** (2004). New evidence of animal consciousness. *Animal Cognition*, 7(1), 5–18.
- Griffiths, R.R., Richards, W.A., Johnson, M.W., McCann, U.D., and Jesse, R.** (2008). Mystical-type experiences occasioned by psilocybin mediate the attribution of personal meaning and spiritual significance 14 months later. *Journal of psychopharmacology*, 22, 621–632.
- Grimes, J.** (1996). On the failure to detect changes in scenes across saccades. In K. Akins (Ed.), *Perception: Vol 2. Vancouver studies in cognitive science* (pp. 89–110). New York: Oxford University Press.
- Grof, S., and Halifax, J.** (1977). *The human encounter with death*. New York: Dutton.
- Gurney, E., Myers, F.W.H., and Podmore, F.** (1886). *Phantasms of the living*. 2 vols. London: Trubner.
- Guterstam, A., and Ehrsson, H.H.** (2012). Disowning one's seen real body during an out-of-body illusion. *Consciousness and Cognition*, 21(2), 1037–1042.
- Hagerty, M. R., Isaacs, J., Brasington, L., Shupe, L., Fetz, E. E., and Cramer, S.C.** (2013). Case study of ecstatic meditation: fMRI and EEG evidence of self-stimulating a reward system. *Neural Plasticity*, article 653572.

• R E F E R E N C E S

Haggard, P. (2008). Human volition: Towards a neuroscience of will. *Nature Reviews Neuroscience*, *9*, 934–946.

Haggard, P. and Clark, S. (2003). Intentional action: Conscious experience and neural prediction. *Consciousness and Cognition*, *12*, 695–707.

Haggard, P., Clark, S., and Kalogeras, J. (2002). Voluntary action and conscious awareness. *Nature Neuroscience*, *5*, 382–385.

Haggard, P., and Eimer, M. (1999). On the relation between brain potentials and the awareness of voluntary movements. *Experimental Brain Research*, *126*, 128–133.

Haggard, P., and Libet, B. (2001). Conscious intention and brain activity. *Journal of Consciousness Studies*, *8*, 47–63.

Haggard, P., Newman, C., and Magno, E. (1999). On the perceived time of voluntary actions. *British Journal of Psychology*, *90*, 291–303.

Hall, C., and Van de Castle, R. (1966). *The content analysis of dreams*. New York: Appleton-Century-Crofts.

Halligan, P., and Oatley, D. (2015). Consciousness isn't all about you, you know. *New Scientist*, 15 August, 26–27.

Hamburger, K., Geremek, A., and Spillmann, L. (2012). Perceptual filling-in of negative coloured afterimages. *Perception*, *41*, 50–56.

Hameroff, S., and Penrose, R. (2014). Consciousness in the universe. A review of the 'Orch OR' theory. *Physics of Life Reviews*, *11*(1), 39–112 (incl. peer commentaries and authors' responses).

Hamilton, W. (1895). Notes and supplementary dissertations. In *Works of Thomas Reid* (pp. 741–1034). 8th ed., vol. 2. Edinburgh: Maclachlan and Stewart.

Hanson, R., and Mendius, R. (2009). *Buddha's brain: The practical science of happiness, love & wisdom*. Oakland, CA: New Harbinger.

Harding, D. E. (1961). *On having no head: Zen and the re-discovery of the obvious*. London: Routledge. Extract reprinted in Hofstadter and Dennett (1981). *The mind's I: Fantasies and reflections on self and soul* (pp. 23–30). London: Penguin.

Hardy, T. (1891). *Tess of the d'Urbervilles: A pure woman faithfully presented*. London: James R. Osgood, McIlvaine & Co. Full text available at <https://www.gutenberg.org/files/110/110-h/110-h.htm> and <https://books.google.co.uk/books?id=nPJAAAAMAAJ>

Hare, B., Call, J., and Tomasello, M. (2001). Do chimpanzees know what conspecifics know? *Animal Behaviour*, *61*(1), 139–151.

Hare, T. A., Camerer, C. F., and Rangel, A. (2009). Self-control in decision-making involves modulation of the vmPFC valuation system. *Science*, *324*, 646–648.

- Harman, G.** (1990). The intrinsic quality of experience. *Philosophical Perspectives*, 4, 31–52.
- Harnad, S.** (1990). The symbol grounding problem. *Physica D*, 42, 335–346.
- Harnad, S.** (2007). Can a machine be conscious? How? *Journal of Consciousness Studies*, 10, 67–75. Also in Clowes, R., Torrance, S., and Chrisley, R. (2007). *Machine consciousness*. Exeter: Imprint Academic.
- Harré, R., and Gillett, G.** (1994). *The discursive mind*. Thousand Oaks, CA: Sage.
- Harris, P.L.** (2000). *The work of the imagination*. Oxford: Blackwell.
- Harris, S.** (2012). *Free will*. New York: Free Press.
- Harris, S.** (2014). *Waking up: A guide to spirituality without religion*. London: Bantam Press.
- Harris, S.** (2014). *The marionette's lament: A response to Daniel Dennett*. www.samharris.org/blog/item/the-marionettes-lament
- Hassin, R.R., Uleman, J.S., and Bargh, J.A.** (Eds) (2005). *The new unconscious*. Oxford: Oxford University Press.
- Hauser, M.D.** (2006). *Moral minds: How nature designed our universal sense of right and wrong*. New York: HarperCollins.
- Hayes, B.** (2015). Computer vision and computer hallucinations. *American Scientist*. www.americanscientist.org/article/computer-vision-and-computer-hallucinations
- Hayhoe, M.** (2000). Vision using routines: A functional account of vision. *Visual Cognition*, 7(1/2/3), 43–64.
- Haynes, J.D.** (2011). Decoding and predicting intentions. *Annals of the New York Academy of Sciences*, 1224, 9–21.
- Hearne, K.** (1978). Lucid dreams: An electrophysiological and psychological study. PhD thesis, University of Hull.
- Hearne, K.** (1990). *The dream machine*. Northants: Aquarian.
- Hebb, D.O.** (1949). *The organization of behaviour: A psychological theory*. New York: Wiley.
- Hermans, H.J.M.** (2011). The dialogical self: A process of positioning in space and time. In S. Gallagher (Ed.), *The Oxford handbook of the self* (pp. 654–680). New York: Oxford University Press.
- Herz, S., and Schooler, J.W.** (2002). A naturalistic study of autobiographical memories evoked by olfactory and visual cues: Testing the Proustian hypothesis. *The American Journal of Psychology*, 115(1), 21–32.

• R E F E R E N C E S

- Hesse, H.** (1943). *The glass bead game* [*Das Glasperlenspiel*]. Zurich: Fretz & Wasmuth. Full text available at https://archive.org/stream/MagisterLudi-TheGlassBeadGame-HermanHesse/hesseludi_djvu.txt (trans. R and C. Winston)
- Heyes, C.M.** (1998). Theory of mind in nonhuman primates. *Behavioral and Brain Sciences*, 21(1), 101–148 (incl. commentaries and author's response).
- Heyes, C.** (2017). Apes submentalise. *Trends in Cognitive Sciences*, 21(1), 1–2.
- Heyes, C.M., and Galef, B.G.** (Eds) (1996). *Social learning in animals: The roots of culture*. San Diego, CA: Academic Press.
- Hilgard, E.R.** (1977). *Divided consciousness: Multiple controls in human thought and action*. New York: Wiley.
- Hirata, S., Watanabe, K., and Kawai, M.** (2001). 'Sweet-potato washing' revisited. In T. Matsuzawa (Ed.), *Primate origins of human cognition and behaviour* (pp. 487–508). Tokyo: Springer.
- Hobbes, T.** (1648/1946). *Leviathan*. Ed. M. Oakeshott. Oxford: Oxford University Press.
- Hobson, J.A.** (1999). *Dreaming as delirium: How the brain goes out of its mind*. Cambridge, MA: MIT Press.
- Hobson, J.A.** (2002). *Dreaming: An introduction to the science of sleep*. New York: Oxford University Press.
- Hobson, J.A.** (2007). Normal and abnormal states of consciousness. In M. Velmans and S. Schneider (Eds), *The Blackwell companion to consciousness* (pp. 101–113). Oxford: Blackwell.
- Hobson, J.A.** (2009). REM sleep and dreaming: Towards a theory of proto-consciousness. *Nature Reviews Neuroscience*, 10, 803–813.
- Hobson, J.A.** (2014). Introduction. In N. Tranquillo (Ed.), *Dream consciousness: Allan Hobson's new approach to the brain and its mind* (pp. 3–8). Cham: Springer International.
- Hobson, J.A., and Friston, K.J.** (2012). Waking and dreaming consciousness: neurobiological and functional considerations. *Progress in neurobiology*, 98(1), 82–98.
- Hobson, J.A., and Friston, K.J.** (2014). Consciousness, dreams, and inference: The Cartesian theatre revisited. *Journal of Consciousness Studies*, 21(1–2), 6–32.
- Hodgson, R.** (1891). A case of double consciousness. *Proceedings of the Society for Psychical Research*, 7, 221–258.
- Hodgson, S.H.** (1870). *The theory of practice: An ethical inquiry: in two books*. London: Longmans, Green, Reader, and Dyer.
- Hofmann, A.** (1980). *LSD, my problem child: Reflections on sacred drugs, mysticism, and science*. New York: McGraw-Hill.

- Hofstadter, D.R.** (1979). *Gödel, Escher, Bach: An eternal golden braid*. New York: Penguin.
- Hofstadter, D.R.** (2007). *I am a strange loop*. London: Penguin.
- Hofstadter, D.R., and Dennett, D.C.** (Eds) (1981). *The mind's I: Fantasies and reflections on self and soul*. London: Penguin.
- Holland, J.** (Ed.) (2001). *Ecstasy: The complete guide: A comprehensive look at the risks and benefits of MDMA*. Rochester, VT: Park Street Press.
- Holland, O.** (2007). A strongly embodied approach to machine consciousness. *Journal of Consciousness Studies*, 14, 97–110.
- Holland, O., and Goodman, R.** (2003). Robots with internal models: A route to machine consciousness? *Journal of Consciousness Studies*, 10(4–5), 77–109.
- Holland, O., Knight, R., and Newcombe, R.** (2007). A robot-based approach to machine consciousness. In A. Chella and R. Manzotti (Eds), *Artificial consciousness* (pp. 156–173). Exeter: Imprint Academic.
- Holmes, D.S.** (1987). The influence of meditation versus rest on physiological arousal. In M. West (Ed.), *The psychology of meditation* (pp. 81–103). Oxford: Clarendon Press.
- Holt, J.** (1999). Blindsight in debates about qualia. *Journal of Consciousness Studies*, 6(5), 54–71.
- Holzinger, B., LaBerge, S., and Levitan, L.** (2006). Psychophysiological correlates of lucid dreaming. *Dreaming*, 16, 88–95.
- Homer.** (8th–7th century BC/1951). *The Iliad of Homer*. Trans. R. Lattimore. Chicago, IL: University of Chicago Press.
- Hong, C.C.-H., Harris, J.C., Pearlson, G.D., Kim, J.S., Calhoun, V.D., Fallon, J.H., . . . and Pekar, J.J.** (2009). fMRI evidence for multisensory recruitment associated with rapid eye movements during sleep. *Human Brain Mapping*, 30(5), 1705–1722.
- Honorton, C., Berger, R.E., Varvoglis, M.P., Quant, M., Derr, P., Schechter, E.I., and Ferrari, D. C.** (1990). Psi communication in the ganzfeld. *Journal of Parapsychology*, 54, 99–139.
- Hood, B., Gjersoe, N.L., and Bloom, P.** (2012). Do children think that duplicating the body also duplicates the mind? *Cognition*, 125, 466–474.
- Hopkins, G.M.** (1885). No worst, there is none. Full text available at <https://www.poetryfoundation.org/poems/44398/no-worst-there-is-none-pitched-past-pitch-of-grief>
- Horikawa, T., Tamaki, M., Miyawaki, Y., and Kamitani, Y.** (2013). Neural decoding of visual imagery during sleep. *Science*, 340(6132), 639–642.
- Horne, J.** (2006). *Sleepfaring: A journey through the science of sleep*. Oxford: Oxford University Press.

• R E F E R E N C E S

- Horne, J.** (2009). Interview with Jon Sutton. *The Psychologist*, 8, 706–709. <http://psychedelicfrontier.com/moving-sacred-world-dmt-nick-sand/>
- Hubbard, B.M.** (1997). *Conscious evolution: Awakening the power of our social potential*. Novato, CA: New World Library.
- Hufford, D.J.** (1982). *The terror that comes in the night: An experience centered study of supernatural assault traditions*. Philadelphia: University of Pennsylvania Press.
- Huijbers, W., Pennartz, C.M., Beldzik, E., Domagalik, A., Vinck, M., Hofman, W.F., Cabeza, R., and Daselaar, S.M.** (2014). Respiration phase-locks to fast stimulus presentations: Implications for the interpretation of posterior midline ‘deactivations’. *Human Brain Mapping*, 35(9), 4932–4943.
- Hume, D.** (1739/2014). *A treatise of human nature*, ed. D.F. Norton and M.J. Norton. Oxford: Oxford University Press (The Clarendon Edition of the Works of David Hume).
- Humphrey, N.** (1983). *Consciousness regained: Chapters in the development of mind*. Oxford: Oxford University Press.
- Humphrey, N.** (1986). *The inner eye: Social intelligence in evolution*. London: Faber & Faber.
- Humphrey, N.** (1987). The inner eye of consciousness. In C. Blakemore and S. Greenfield (Eds), *Mindwaves: Thoughts on intelligence, identity, and consciousness* (pp. 377–381). Oxford: Blackwell.
- Humphrey, N.** (1992). *A history of the mind: Evolution and the birth of consciousness*. London: Chatto & Windus.
- Humphrey, N.** (2000). How to solve the mind–body problem. *Journal of Consciousness Studies*, 7(4), 5–20, with commentaries, pp. 21–97, and reply, pp. 98–112. Reprinted in Humphrey (2002). *The mind made flesh: Frontiers of psychology and evolution* (pp. 90–114). Oxford: Oxford University Press.
- Humphrey, N.** (2002). *The mind made flesh: Essays from the frontiers of psychology and evolution*. Oxford: Oxford University Press.
- Humphrey, N.** (2006). *Seeing red: A study in consciousness*. Cambridge, MA: Harvard University Press.
- Humphrey, N.** (2011). *Soul dust: The magic of consciousness*. London: Quercus.
- Humphrey, N.** (2016). Redder than red illusionism or phenomenal surrealism? *Journal of Consciousness Studies*, 23(11–12), 116–123.
- Humphrey, N.** (2017). The invention of consciousness. *Topoi*. <https://doi.org/10.1007/s11245-017-9498-0>
- Humphrey, N., and Dennett, D.C.** (1989). Speaking for our selves: An assessment of multiple personality disorder. *Raritan*, 9(1), 68–98.

- Humphreys, C.** (1951). *Buddhism: An introduction and guide*. Harmondsworth: Penguin.
- Hunt, A. R., and Kingstone, A.** (2003). Covert and overt voluntary attention: Linked or independent? *Cognitive Brain Research*, 18, 102–105.
- Hunt, H.** (2006). The truth value of mystical experiences. *Journal of Consciousness Studies*, 13(12), 5–43.
- Hurley, S.L.** (1998). *Consciousness in action*. Cambridge, MA: Harvard University Press.
- Hurley, S.L.** (2001). Perception and action: Alternative views. *Synthese*, 129(1), 3–40.
- Hut, P.** (1999). Theory and experiment in philosophy. *Journal of Consciousness Studies*, 6(2–3), 241–244. Reprinted in F.J. Varela and J. Shear (Eds) (1999). *The view from within* (pp. 241–244). Thorverton, Devon: Imprint Academic.
- Huxley, A.** (1954). *The doors of perception*. London: Chatto & Windus.
- Hyman, I.E., Boss, S.M., Wise, B.M., and Caggiano, J.M.** (2010). Did you see the unicycling clown? Inattentive blindness while walking and talking on a cell phone. *Applied Cognitive Psychology*, 24(5), 597–607.
- Hyman, R.** (1985). The ganzfeld psi experiment: A critical appraisal. *Journal of Parapsychology*, 49, 3–49.
- Hyman, R.** (1995). Evaluation of the program on anomalous mental phenomena. *Journal of Parapsychology*, 59, 321–351.
- Hyman, R., and Honorton, C.** (1986). A joint communique: The psi ganzfeld controversy. *Journal of Parapsychology*, 50, 351–364.
- International Association for the Study of Pain (IASP).** (2011). Part III: Pain terms: A current list with definitions and notes on usage. In *Classification of chronic pain*. 2nd ed. <https://www.iasp-pain.org/files/Content/ContentFolders/Publications2/FreeBooks/Classification-of-Chronic-Pain.pdf>
- Inugami, M., and Ma, T.I.M.** (2002). Factors related to the occurrence of isolated sleep paralysis elicited during a multi-phasic sleep-wake schedule. *Sleep*, 25(1), 89.
- Ionta, S., Heydrich, L., Lenggenhager, B., Mouthon, M., Fornari, E., Chapuis, D., . . . and Blanke, O.** (2011). Multisensory mechanisms in temporo-parietal cortex support self-location and first-person perspective. *Neuron*, 70(2), 363–374.
- Irwin, D.E.** (1991). Information integration across saccadic eye movements. *Cognitive Psychology*, 23, 420–456.
- Irwin, H. J.** (1985). *Flight of mind: A psychological study of the out-of-body experience*. Metuchen, NJ: Scarecrow Press.

• R E F E R E N C E S

- Irwin, H. J., and Watt, C. A.** (2007). *An introduction to parapsychology*. 5th ed. Jefferson, NC: McFarland.
- Jablonska, E., Lamb, M. J., and Zeligowski, A.** (2005). *Evolution in four dimensions: Genetic, epigenetic, behavioral and symbolic variation in the history of life*. Cambridge, MA: MIT Press.
- Jackendoff, R.** (1987). *Consciousness and the computational mind*. Cambridge, MA: MIT Press.
- Jackson, F.** (1982). Epiphenomenal qualia. *Philosophical Quarterly*, 32, 127–136. Reprinted in Chalmers, D. (2002). *Philosophy of mind: Classical and contemporary readings* (pp. 273–280). New York: Oxford University Press.
- Jackson, F.** (1998). Postscript on qualia. In F. Jackson, *Mind, methods and conditionals: Selected papers* (pp. 76–79). London: Routledge.
- Jackson, F.** (2003). Mind and illusion. *Royal Institute of Philosophy Supplement*, 53, 251–271.
- Jajdelska, E., Butler, C., Kelly, S., McNeill, A., and Overy, K.** (2011). Crying, moving, and keeping it whole: What makes literary description vivid? *Poetics Today*, 31(3), 433–463.
- James, H.** (1881). *The portrait of a lady*. Boston, MA: Houghton, Mifflin and Co.
- James, W.** (1890). *The principles of psychology*. 2 vols. London: MacMillan.
- James, W.** (1902). *The varieties of religious experience: A study in human nature*. New York: Longmans, Green and Co.
- James, W.** (1904). Does consciousness exist? *Journal of Philosophy, Psychology, and Scientific Methods*, 1, 477–491.
- James, W.** (1907/1975). *Pragmatism: A new name for some old ways of thinking*. Cambridge, MA: Harvard University Press.
- Jamieson, G. A.** (2005). The modified Tellegen Scale: A clearer window on the structure and meaning of absorption. *Australian Journal of Clinical and Experimental Hypnosis*, 33(2), 119–139.
- Jansen, K.** (2001). *Ketamine: Dreams and realities*. Sarasota, FL: Multidisciplinary Association for Psychedelic Studies.
- Jay, M.** (2009). *The atmosphere of heaven: The unnatural experiments of Dr. Beddoes and his sons of genius*. New Haven, CT: Yale University Press.
- Jaynes, J.** (1976). *The origin of consciousness in the breakdown of the bicameral mind*. New York: Houghton Mifflin.
- Jenner, J. A., and van Laar, T.** (2011). Visual hallucinations in Parkinson's disease. In V.C.N. Wong and V.L.Y. Chu (Eds), *International encyclopedia of rehabilitation*. Center for International Rehabilitation Research Information and Exchange (CIRRIE).

- Johansson, R., Holsanova, J., and Holmqvist, K.** (2006). Pictures and spoken descriptions elicit similar eye movements during mental imagery, both in light and in complete darkness. *Cognitive Science*, 30, 1053–1079.
- Johnson, B. R., and Lam, S. K.** (2010). Self-organization, natural selection, and evolution: Cellular hardware and genetic software. *BioScience*, 60(11), 879–885.
- Johnson, M. K., and Raye, C. L.** (1981). Reality monitoring. *Psychological Review*, 88, 67–85.
- Joyce, J.** (1922). *Ulysses*. Paris: Sylvia Beach. Full text available at <https://www.gutenberg.org/files/4300/4300-h/4300-h.htm>
- Joyce, R.** (2007). *The evolution of morality*. Cambridge, MA: MIT Press.
- Julien, R. M.** (2001). *A primer of drug action: A concise, nontechnical guide to the actions, uses, and side effects of psychoactive drugs*. Rev. ed. New York: Henry Holt.
- Jung, C. G.** (1934–1936/1968). *The archetypes and the collective unconscious [Die Archetypen und das kollektive Unbewußte]*. 2nd ed. Trans. R.F.C. Hull. London: Routledge.
- Jung, R. E., Mead, B. S., Carrasco, J., and Flores, R. A.** (2013). The structure of creative cognition in the human brain. *Frontiers in Human Neuroscience*, 7, article 330.
- Kabadayi, C., and Osvath, M.** (2017). Ravens parallel great apes in flexible planning for tool-use and bartering. *Science*, 357, 202–204.
- Kabat-Zinn, J.** (1999). Indra’s net at work: The mainstreaming of dharma practice in society. In G. Watson, S. Batchelor, and G. Claxton (Eds), *The psychology of awakening: Buddhism, science and our day-to-day lives* (pp. 225–249). London: Rider.
- Kabat-Zinn, J.** (2003). Mindfulness-based interventions in context: Past, present, and future. *Clinical Psychology: Science and Practice*, 10(2), 144–156.
- Kafatos, M., Tanzi, R. E., and Chopra, D.** (2011). How consciousness becomes the physical universe. *Journal of Cosmology*, 14. <http://journalofcosmology.com/Consciousness140.html>
- Kafka, F.** (1990). *Tagebücher*. Ed. H.-G. Koch, M. Müller, and M. Pasley. New York: Schocken.
- Kahneman, D.** (2011). *Thinking, fast and slow*. New York: Farrar, Straus, and Giroux.
- Kallio, S., and Revonsuo, A.** (2003). Hypnotic phenomena and altered states of consciousness: a multilevel framework of description and explanation. *Contemporary Hypnosis*, 20, 111–164. Peer commentaries in *Contemporary Hypnosis* (2005), 22(1), 1–55.

• R E F E R E N C E S

Kaminski, J., Call, J., and Tomasello, M. (2008). Chimpanzees know what others know, but not what they believe. *Cognition*, *109*, 224–234.

Kanwisher, N. (2001). Neural events and perceptual awareness. *Cognition*, *79*, 89–113. Reprinted in S. Dehaene (Ed.) (2002), *The cognitive neuroscience of consciousness* (pp. 89–113). Cambridge, MA: MIT Press.

Kapleau, R.P. (1980). *The three pillars of Zen: Teaching, practice, and enlightenment*. Rev. ed. New York: Doubleday.

Karn, K. and Hayhoe, M. (2000). Memory representations guide targeting eye movements in a natural task. *Visual Cognition*, *7*, 673–703.

Karremans, J.C., Stroebe, W., and Claus, J. (2006). Beyond Vicary's fantasies: The impact of subliminal priming and brand choice. *Journal of Experimental Social Psychology*, *42*, 792–798.

Kasamatsu, A., and Hirai, T. (1966). An electroencephalographic study on the Zen meditation (zazen) *Folia Psychiatrica et Neurologica Japonica*, *20*, 315–336.

Kasparov, G. (2017). *Deep thinking: Where machine intelligence ends and human creativity begins*. London: Hachette.

Kassewitz, J., Hyson, M.T., Reid, J.S., and Barrera, R.L. (2016). A phenomenon discovered while imaging dolphin echolocation sounds. *Marine Science: Research & Development*, *6*(4), article 1000202.

Kathirvel, N., and Mortimer, A. (2013). Causes, diagnosis and treatment of visceral hallucinations. *Progress in Neurology and Psychiatry*, Jan/Feb, 6–10.

Katz, S.T. (1978). Language, epistemology, and mysticism. In S.T. Katz (Ed.), *Mysticism and philosophical analysis* (pp. 22–74). New York: Oxford University Press.

Kędzierski, J., Muszyński, R., Zoll, C., Oleksy, A., and Frontkiewicz, M. (2013). EMYS – emotive head of a social robot. *International Journal of Social Robotics*, *5*(2), 237–249.

Keeley, B. (2009). Early history of the quale and its relation to the senses. *Routledge Companion to Philosophy of Psychology*, ed. J. Symons and P. Calvo (pp. 71–89). London: Routledge.

Kemmerer, D. (2015). Are we ever aware of concepts? A critical question for the global neuronal workspace, integrated information, and attended intermediate-level representation theories of consciousness. *Neuroscience of Consciousness*, *1*, niv006.

Kentridge, R.W., and Heywood, C.A. (1999). The status of blindsight: Near-threshold vision, islands of cortex and the Riddoch phenomenon. *Journal of Consciousness Studies*, *6*(5), 3–11.

- Kessler, K., and Braithwaite, J.** (2016). Deliberate and spontaneous sensations of disembodiment: Capacity or flaw? *Neuropsychiatry*, 21(5), 412–428.
- Key, B.** (2016). Why fish do not feel pain. *Animal Sentience: An Interdisciplinary Journal on Animal Feeling*, 1(3), 39.
- Khoury, B., Lecomte, T., Fortin, G., Masse, M., Therien, P., Bouchard, V., Chapleau, M.-A., Paquin, K., and Hofmann, S. G.** (2013). Mindfulness-based therapy: A comprehensive meta-analysis. *Clinical Psychology Review*, 33, 763–771.
- Kihlstrom, J. F.** (1985). Hypnosis. *Annual Review of Psychology*, 36, 385–418.
- Kihlstrom, J. F.** (1987). The cognitive unconscious. *Science*, 237, 1445–1638.
- Kihlstrom, J. F.** (1996). Perception without awareness of what is perceived, learning without awareness of what is learned. In M. Velmans (Ed.), *The science of consciousness* (pp. 23–46). London: Routledge.
- Kihlstrom, J. F., and Cork, R. C.** (2007). Consciousness and anesthesia. In S. Schneider and M. Velmans (Eds), *The Blackwell companion to consciousness* (pp. 628–639). Chichester: Wiley.
- Kim, J.** (2007). The causal efficacy of consciousness. In M. Velmans and S. Schneider (Eds), *The Blackwell companion to consciousness* (pp. 406–417). Oxford: Blackwell.
- Kirk, R.** (2005). *Zombies and consciousness*. Oxford: Clarendon.
- Kirk, R.** (2015). Zombies. In E. N. Zalta (Ed.), *The Stanford encyclopedia of philosophy*, Summer 2015 Edition. <http://plato.stanford.edu/archives/sum2015/entries/zombies/>
- Kirk, R., and Squires, J.E.R.** (1974). Zombies vs materialists. *Proceedings of the Aristotelian Society*, 48, 135–152.
- Kirsch, I.** (1997). The state of the altered state debate. In W. J. Matthews and J. H. Edgette (Eds), *Current thinking and research in brief therapy: Solutions, strategies, narratives* (pp. 91–99). New York: Brunner/Mazel.
- Kirsch, I.** (2011). The altered state issue: Dead or alive? *International Journal of Clinical and Experimental Hypnosis*, 59(3), 350–362.
- Klein, C., and Barron, A. B.** (2016). Insects have the capacity for subjective experience. *Animal Sentience*, 9(1).
- Klüver, H.** (1926). Mescal visions and eidetic vision. *American Journal of Psychology*, 37, 502–515.
- Kobes, B. W.** (2007). Functionalist theories of consciousness. In T. Bayne, A. Cleeremans, and P. Wilken (Eds), *Oxford companion to consciousness* (pp. 310–315). Oxford: Oxford University Press.

• R E F E R E N C E S

- Koch, C.** (2004). *The quest for consciousness: A neurobiological approach*. Englewood, CO: Roberts & Co.
- Koch, C., and Hepp, K.** (2006/2007). Quantum mechanics in the brain. *Nature*, 440, 611–612. <https://authors.library.caltech.edu/40450/7/561.pdf>
- Koch, C., Massimini, M., Boly, M., and Tononi, G.** (2016). Neural correlates of consciousness: Progress and problems. *Nature Reviews Neuroscience*, 17, 307–321.
- Koch, C., and Tsuchiya, N.** (2007). Attention and consciousness: Two distinct brain processes. *Trends in Cognitive Sciences*, 11(1), 16–22.
- Komatsu, H.** (2006). The neural mechanisms of perceptual filling-in. *Nature Reviews: Neuroscience*, 7, 220–231.
- Kosslyn, S. M., Ball, T.M., and Reiser, B.J.** (1978). Visual images preserve metric spatial information: Evidence from studies of image scanning. *Journal of Experimental Psychology: Human Perception and Performance*, 4(1), 47–60.
- Kosslyn, S.M.** (1980). *Image and mind*. Cambridge, MA: Harvard University Press.
- Kosslyn, S.M., Thompson, W.L., and Ganis, G.** (2006). *The case for mental imagery*. Oxford: Oxford University Press.
- Kozuch, B.** (2015). The received method for ruling out brain areas from being NCC undermines itself. *Journal of Consciousness Studies*, 22(9–10), 145–169.
- Krippner, S.** (2000). The epistemology and technologies of shamanic states of consciousness. *Journal of Consciousness Studies*, 7(11–12), 93–118.
- Krupenye, C., Kano, F., Hirata, S., Call, J., and Tomasello, M.** (2016). Great apes anticipate that other individuals will act according to false beliefs. *Science*, 354(6308), 110–114.
- Kryger, M.H., Roth, T.R., and Dement, W.C.** (2011). *Principles and practice of sleep medicine*. 5th ed. St Louis: Elsevier.
- Kuhn, G., Amlani, A.A., and Rensink, R.A.** (2008). Towards a science of magic. *Trends in Cognitive Sciences*, 12, 349–353.
- Kuhn, T.S.** (1962). *The structure of scientific revolutions*. Chicago: University of Chicago Press.
- Kunst-Wilson, W.R., and Zajonc, R.B.** (1980). Affective discrimination of stimuli that cannot be recognized. *Science*, 207, 557–558.
- Kurzweil, R.** (1999). *The age of spiritual machines: When computers exceed human intelligence*. New York: Texere.
- Kurzweil, R.** (2005). *The singularity is near: When humans transcend biology*. New York: Viking.

- Kuwamura, K., Nishio, S., and Sato, S.** (2016). Can we talk through a robot as if face-to-face? Long-term fieldwork using teleoperated robot for seniors with Alzheimer's disease. *Frontiers in Psychology*, 7, article 1066.
- LaBerge, S.** (1985). *Lucid dreaming*. Los Angeles: Tarcher.
- LaBerge, S.** (1988). The psychophysiology of lucid dreaming. In J. Gackenbach (Ed.), *Conscious mind, sleeping brain* (pp. 135–153). New York: Springer.
- LaBerge, S.** (1990). Lucid dreaming: Psychophysiological studies of consciousness during REM sleep. In R. R. Bootzen, J. F. Kihlstrom, and D. L. Schacter (Eds), *Sleep and cognition* (pp. 109–126). Washington, DC: American Psychological Association.
- LaBerge, S.** (2000). Lucid dreaming: Evidence and methodology. *Behavioral and Brain Sciences*, 23(6), 962–963. Commentary on J. Hobson, E. F. Pace-Schott, and R. Stickgold, 'Dreaming and the brain: Toward a cognitive neuroscience of conscious states', *BBS*, 23(6), 793–1035. Version with figures added at www.lucidity.com/slbbbs/index.html
- LaBerge, S., and Rheingold, H.** (1990). *Exploring the world of lucid dreaming*. New York: Ballantine Books.
- Lakoff, G., and Johnson, M.** (1980/2003). *Metaphors we live by*. Chicago: University of Chicago Press.
- Lakoff, G., and Johnson, M.** (1999). *Philosophy in the flesh: The embodied mind and its challenge to Western thought*. New York: Basic Books.
- Laloyaux, C., Devue, C., Doyen, S., David, E. and Cleeremans, A.** (2008). Undetected changes in visible stimuli influence subsequent decisions. *Consciousness and Cognition*, 17, 646–656.
- Lama, D., Benson, H., Thurman, R., Gardner, H., and Goleman, D.** (1991). *MindScience: An east-west dialogue*. Somerville, MA: Mind/Body Medical Institute Inc. & Tibet House New York Inc.
- La Mettrie, J. O. de** (1748 anon./2015). *L'homme machine [Man – machine]*. Trans. J. Bennett. www.earlymoderntexts.com/assets/pdfs/lamettrie1748.pdf
- Lanier, J.** (1995). You can't argue with a zombie. *Journal of Consciousness Studies*, 2(4), 333–345.
- Lau, H. C., Rogers, R. D., Ramnani, N., and Passingham, R. E.** (2004a). Willed action and attention to the selection of action. *NeuroImage*, 21, 1407–1415.
- Lau, H. C., Rogers, R. D., Haggard, P., and Passingham, R. E.** (2004b). Attention to intention. *Science*, 303(5661), 1208–1210.
- Laureys, S.** (2005). The neural correlate of (un)awareness: Lessons from the vegetative state. *Trends in Cognitive Sciences*, 9, 556–559.

• R E F E R E N C E S

Laureys, S. (2009). Arousal vs awareness. In T. Bayne, A. Cleeremans, and P. Wilken (Eds), *The Oxford companion to consciousness* (pp. 58–60). Oxford: Oxford University Press.

Lavazza, A., and Robinson, H. (Eds). (2014). *Contemporary dualism: A defense*. New York: Routledge.

Lavie, N., Beck, D. M., and Konstantinou, N. (2014). Blinded by the load: Attention, awareness and the role of perceptual load. *Philosophical Transactions of the Royal Society B*, 369, 20130205.

Lawrence, D. H. (1913). *Sons and lovers*. London: Gerald Duckworth and Company. Full text available at <https://www.gutenberg.org/files/217/217-h/217-h.htm>

Leary, T. (1983). *Flashbacks: A personal and cultural history of an era: an autobiography*. London: Heinemann, and Los Angeles: Jeremy P. Tarcher.

Lebedev, A. V., Lövdén, M., Rosenthal, G., Feilding, A., Nutt, D. J., and Carhart-Harris, R. L. (2015). Finding the self by losing the self: Neural correlates of ego-dissolution under psilocybin. *Human Brain Mapping*, 36(8), 3137–3153.

Lebedev, A. V., Kaelen, M., Lövdén, M., Nilsson, J., Feilding, A., Nutt, D. J., and Carhart-Harris, R. L. (2016). LSD-induced entropic brain activity predicts subsequent personality change. *Human Brain Mapping*, 37(9), 3203–3213.

Leibniz, G. W. (1714/1898). *Monadology and other philosophical writings [La Monadologie]*. Trans. R. Latta. Oxford: Oxford University Press.

Lem, S. (1981). The seventh sally or How Trurl's own perfection led to no good. In D. R. Hofstadter and D. C. Dennett (Eds), *The mind's I* (pp. 287–295, with commentary). London: Penguin.

Lenggenhager, B., Arnold, C. A., and Giummarra, M. J. (2014). Phantom limbs: Pain, embodiment, and scientific advances in integrative therapies. *Wiley Interdisciplinary Reviews: Cognitive Science*, 5(2), 221–231.

Lenggenhager, B., Mouthon, M., and Blanke, O. (2009). Spatial aspects of bodily self-consciousness. *Consciousness and Cognition*, 18(1), 110–117.

Lenggenhager, B., Tadi, T., Metzinger, T., and Blanke, O. (2007). Video ergo sum: Manipulating bodily self-consciousness. *Science*, 317(5841), 1096–1099.

Lettvin, J. Y., Maturana, H. R., McCulloch, W. S., and Pitts, W. H. (1968). What the frog's eye tells the frog's brain. *Proceedings of the Institute of Radio Engineers*, 47, 1940–1951.

Levin, D. T. (2002). Change blindness blindness: as visual metacognition. *Journal of Consciousness Studies*, 9(5–6), 111–130.

Levine, J. (2001). *Purple haze: The puzzle of consciousness*. New York: Oxford University Press.

Levine, S. (1979). *A gradual awakening*. New York: Doubleday.

Levinson, B.W. (1965). States of awareness during general anaesthesia: Preliminary communication. *British Journal of Anaesthesia*, 37, 544–546.

Libet, B., Wright, E.W. Jr., Feinstein, B., and Pearl, D.K. (1979). Subjective referral of the timing for a conscious sensory experience. *Brain*, 102, 193–224.

Libet, B. (1982). Brain stimulation in the study of neuronal functions for conscious sensory experiences. *Human Neurobiology*, 1, 235–242.

Libet, B. (1985). Unconscious cerebral initiative and the role of conscious will in voluntary action. *Behavioral and Brain Sciences*, 8(4), 529–566 (incl. commentaries and author's response; discussion continued in *BBS*, 10(2), 318–321).

Libet, B. (1991). Scientific correspondence. *Nature*, 351, 195.

Libet, B. (1999). Do we have free will? *Journal of Consciousness Studies*, 6(8–9), 47–57. Also reprinted in B. Libet, A. Freeman, and K. Sutherland (Eds) (1999). *The volitional brain: Towards a neuroscience of free will* (pp. 47–57). Thorverton, Devon: Imprint Academic.

Libet, B. (2004). *Mind time: The temporal factor in consciousness*. Cambridge, MA: Harvard University Press.

Libet, B., Gleason, C.A., Wright, E.W., and Pearl, D.K. (1983). Time of conscious intention to act in relation to onset of cerebral activity (readiness potential): The unconscious initiation of a freely voluntary act. *Brain*, 106, 623–642.

Libet, B., Wright, E.W., Feinstein, B., and Pearl, D.K. (1979). Subjective referral of the timing for a conscious sensory experience: a functional role for the somatosensory specific projection system in man. *Brain*, 102, 191–222.

Limanowski, J. (2014). What can body ownership illusions tell us about minimal phenomenal selfhood? *Frontiers in Human Neuroscience*, 8, 946.

Limanowski, J., and Hecht, H. (2011). Where do we stand on locating the self? *Psychology*, 2(4), 312–317.

Lin, P., Abney, K., and Bekey, G.A. (2011). *Robot ethics: The ethical and social implications of robotics*. Cambridge, MA: MIT Press.

Lindahl, J.R., Kaplan, C.T., Winget, E.M., and Britton, W.B. (2014). A phenomenology of meditation-induced light experiences: Traditional Buddhist and neurobiological perspectives. *Frontiers in Psychology*, 4, article 973.

• R E F E R E N C E S

Lindberg, D.C. (1976). *Theories of vision from Al-Kindi to Kepler*. Chicago: University of Chicago Press.

Linzey, A. (2009). *Why animal suffering matters: Philosophy, theology, and practical ethics*. New York: Oxford University Press.

Lippelt, D.P., Hommel, B., and Colzato, L.S. (2014). Focused attention, open monitoring and loving kindness meditation: Effects on attention, conflict monitoring, and creativity – a review. *Frontiers in Psychology*, 5, article 1083.

Llinás, R.R. (2002). *I of the Vortex: From neurons to self*. Cambridge, MA: MIT Press.

Llinás, R., and Ribary, U. (2001). Consciousness and the brain. *Annals of the New York Academy of Sciences*, 929(1), 166–175.

Lloyd, D. (2000). Popping the thought balloon. In D. Ross, A. Brook, and D. Thompson (Eds), *Dennett's philosophy: A comprehensive assessment* (p. 169–199). Cambridge, MA: MIT Press.

Lloyd, D. (2004). *Radiant cool: A novel theory of consciousness*. Cambridge, MA: MIT Press.

Lodge, D. (2001). *Thinks. . . .* London: Secker & Warburg.

Loftus, E., and Ketcham, K. (1994). *The myth of repressed memory: False memories and allegations of sexual abuse*. New York: St Martin's Press.

Loftus, E.F., and Palmer, J.C. (1974). Reconstruction of auto-mobile destruction: An example of the interaction between language and memory. *Journal of Verbal Learning and Verbal Behavior*, 13, 585–589.

Logothetis, N.K., and Schall, J.D. (1989). Neuronal correlates of subjective visual perception. *Science*, 245, 761–763.

Longo, M.R., and Tsakiris, M. (2013). Merging second-person and first-person neuroscience. *Behavioral and Brain Sciences*, 36(4), 429–430. Commentary on Schilbach et al., 'Toward a second-person neuroscience?', *BBS*, 36(4), 393–462.

Louie, K., and Wilson, M. (2001). Temporally structured replay of awake hippocampal ensemble activity during rapid eye movement sleep. *Neuron*, 29, 145–156.

Ludlow, P., Nagasawa, Y., and Stoljar, D. (Eds) (2004). *There's something about Mary: Essays on phenomenal consciousness and Frank Jackson's knowledge argument*. Cambridge, MA: MIT Press.

Ludwig, K. (2002). The mind-body problem: An overview. In S.P. Stich and T.A. Warfield (Eds), *The Blackwell guide to philosophy of mind* (pp. 1–46). Malden, MA: Blackwell.

- Lumer, E. D.** (2000). Binocular rivalry and human visual awareness. In T. Metzinger (Ed.), *Neural correlates of consciousness* (pp. 231–240). Cambridge, MA: MIT Press.
- Lumer, E. D., Friston, K. J., and Rees, G.** (1998). Neural correlates of perceptual rivalry in the human brain. *Science*, *280*, 1930–1934.
- Lumma, A. L., Koko, B. E., and Singer, T.** (2015). Is meditation always relaxing? Investigating heart rate, heart rate variability, and likeability during training of three types of meditation. *International Journal of Psychophysiology*, *97*(1), 38–45.
- Luna, L. E.** (2016). Some observations on the phenomenology of the ayahuasca experience. In L. E. Luna and S. F. White (Eds), *Ayahuasca reader: Encounters with the Amazon's sacred vine* (pp. 251–279). Santa Fe: Synergetic Press.
- Luna, L. E., and White, S. F.** (2016). *Ayahuasca reader: Encounters with the Amazon's sacred vine*. Santa Fe: Synergetic Press.
- Luria, A. R.** (1968). *The mind of a mnemonist: A little book about a vast memory*. Trans. L. Solotaroff. London: Jonathan Cape.
- Lutz, A., Lachaux, J.-P., Martinerie, J., and Varela, F. J.** (2002). Guiding the study of brain dynamics by using first-person data: Synchrony patterns correlate with conscious states during a simple visual task. *Proceedings of the National Academy of Sciences, USA*, *99*(3), 1586–1591.
- Lutz, A., Slagter, H. A., Dunne, J. D., and Davidson, R. J.** (2008). Attention regulation and monitoring in meditation. *Trends in Cognitive Sciences*, *12*(4), 163–169.
- Lutz, J.** (2016). Neural correlates of mindfulness: Investigating self-related processes in mindfulness meditators using functional magnetic resonance imaging. PhD diss, University of Zurich.
- Lycan, W. G.** (2004). The superiority of HOP to HOT. In R. Gennaro (Ed.), *Higher-order theories of consciousness: An anthology* (pp. 93–114). Amsterdam: John Benjamins.
- Lynn, M. T., Muhle-Karbe, P. S., Aarts, H., and Brass, M.** (2014). Priming determinist beliefs diminishes implicit (but not explicit) components of self-agency. *Frontiers in Psychology*, *5*, article 1483.
- MacIntyre, A.** (1985). *After virtue*. 2nd ed. Notre Dame, IN: University of Notre Dame Press.
- Mack, A.** (2003). Inattention blindness: Looking without seeing. *Current Directions in Psychological Science*, *12*, 180–184.
- Mack, A., and Rock, I.** (1998). *Inattention blindness*. Cambridge, MA: MIT Press.
- MacKay, D.** (1987). Divided brains – divided minds? In C. Blakemore and S. Greenfield (Eds), *Mindwaves* (pp. 5–16). Oxford: Blackwell.

• R E F E R E N C E S

- Macknik, S.L., King, M., Randi, J., Robbins, A. Teller, J.T., and Martinez-Conde, S.** (2008). Attention and awareness in stage magic: Turning tricks into research. *Nature Reviews Neuroscience*, 9, 871–879.
- Macphail, E.M.** (1998). *The evolution of consciousness*. Oxford: Oxford University Press.
- Mahr, J., and Csibra, G.** (2017). Why do we remember? The communicative function of episodic memory. *Behavioral and Brain Sciences*, 1–93.
- Maia, T.V., and Cleeremans, A.** (2005). Consciousness: Converging insights from connectionist modeling and neuroscience. *Trends in Cognitive Sciences*, 9(8), 397–404.
- Malcolm, N.** (1959). *Dreaming*. London: Routledge and Kegan Paul.
- Malinowski, P.** (2013). Neural mechanisms of attentional control in mindfulness meditation. *Frontiers in Neuroscience*, 7, article 8.
- Maloney, J.C.** (1985). About being a bat. *Australasian Journal of Philosophy*, 63, 26–49.
- Mandler, G.** (2007). Involuntary memories: Variations on the unexpected. In J.H. Mace (Ed.), *Involuntary memory* (pp. 50–67). Malden, MA: Blackwell.
- Mangan, B.** (2001). Sensation's ghost: The non-sensory 'fringe' of consciousness. *Psyche*, 7(18). <https://www.semanticscholar.org/paper/Sensation-s-Ghost-The-Non-Sensory-Fringe-of-Consci-Mangan/d0e1c815febb853317f86147f25d5d1d7a41c8c1>
- Mann, T.** (1912). *Death in Venice [Der Tod in Venedig]*. Berlin: S. Fischer.
- Manuello, J., Vercelli, U., Nani, A., Costa, T., and Cauda, F.** (2016). Mindfulness meditation and consciousness: An integrative neuroscientific perspective. *Consciousness and Cognition*, 40, 67–78.
- Maquet, P., Ruby, P., Maudoux, A., Albouy, G., Sterpenich, V., Dang-Vu, T., . . . and Laureys, S.** (2005). Human cognition during REM sleep and the activity profile within frontal and parietal cortices: A reappraisal of functional neuroimaging data. *Progress in Brain Research*, 150, 219–227.
- Marcel, A.J.** (1983). Conscious and unconscious perception: Experiments on visual masking and word recognition. *Cognitive Psychology*, 15, 197–237.
- Marks, D.** (2000). *The psychology of the psychic*. 2nd ed. Buffalo, NY: Prometheus.
- Marques, H.G., and Holland, O.** (2009). Architectures for functional imagination. *Neurocomputing*, 72, 743–759.
- Marshall, J., and Halligan, P.** (1988). Blindsight and insight in visuo-spatial neglect. *Nature*, 336, 766–777.
- Martin, A., and Santos, L.R.** (2016). What cognitive representations support primate theory of mind? *Trends in Cognitive Sciences*, 20(5), 375–382.

- Mason, L., Peters, E., Williams, S.C., and Kumari, V.** (2017). Brain connectivity changes occurring following cognitive behavioural therapy for psychosis predict long-term recovery. *Translational Psychiatry*, 7(1), e1001.
- Masters, R.E. and Houston, J.** (1967). *The varieties of psychedelic experience*. London: Anthony Blond.
- Maughan, P.** (2017). Could ‘microdoses’ of LSD be used to treat depression? *New Statesman*, 16 February. www.newstatesman.com/culture/books/2017/02/could-microdoses-ldd-be-used-help-depression
- Maury, A.** (1861). *Le sommeil et les rêves: Études psychologiques sur ces phénomènes et les divers états qui s’y attachent*. Paris: Didier.
- Mauskopf, S.H., and McVaugh, M.R.** (1980). *The elusive science: Origins of experimental psychical research*. Baltimore: Johns Hopkins University Press.
- Mavromatis, A.** (1987). *Hypnagogia: The unique state of consciousness between wakefulness and sleep*. London: Routledge and Kegan Paul.
- Mazzoni, G., Venneri, A., McGeown, W. J., and Kirsch, I.** (2013). Neuroimaging resolution of the altered state hypothesis. *Cortex*, 49, 400–410.
- McCull, D., and Nejat, G.** (2014). Recognizing emotional body language displayed by a human-like social robot. *International Journal of Social Robotics*, 6(2), 261–280.
- McCreery, C., and Claridge, G.** (2002). Healthy schizotypy: The case of out-of-the-body experiences. *Personality and Individual Differences*, 32(1), 141–154.
- McCrone, J.** (1999). *Going inside: A tour round a single moment of consciousness*. London: Faber & Faber.
- McEwan, I.** (2005). *Saturday*. London: Jonathan Cape.
- McGinn, C.** (1987). Could a machine be conscious? In C. Blakemore and S. Greenfield (Eds), *Mindwaves* (pp. 279–288). Oxford: Blackwell.
- McGinn, C.** (1991). *The problem of consciousness: Essays towards a resolution*. Cambridge, MA: Blackwell.
- McGinn, C.** (1999). *The mysterious flame: Conscious minds in a material world*. New York: Basic Books.
- McGrath, J.J., Saha, S., Al-Hamzawi, A., Alonso, J., Bromet, E.J., Bruffaerts, R., . . . and Kessler, R.C.** (2015). Psychotic experiences in the general population: A cross-national analysis based on 31 261 respondents from 18 countries. *JAMA Psychiatry*, 72(7), 697–705.
- McMains, S.A., and Somers, D.C.** (2004). Multiple spotlights of attentional selection in human visual cortex. *Neuron*, 42, 677–686.

• R E F E R E N C E S

- McNally, R.J.** (2012). Searching for repressed memory. In R. F. Belli (Ed.), *True and false recovered memories: Toward a reconciliation of the debate* (pp. 121–147). New York: Springer.
- Meadi, A., Jennum, P., Ptito, M., and Kupers, R.** (2014). The sensory construction of dreams and nightmare frequency in congenitally blind and late blind individuals. *Sleep Medicine, 15*(5), 586–595.
- Medina, F.S., Taylor, A.H., Hunt, G.R., and Gray, R.D.** (2011). New Caledonian crows' responses to mirrors. *Animal Behaviour, 82*, 981–993.
- Meijer, P.** (2002). *Seeing with sound for the blind: Is it vision?* Paper presented at Toward a Science of Consciousness, Tucson, AZ, April 8–12, 2002. Conference Research Abstracts (provided by *Journal of Consciousness Studies*), Abstract No. 187. Soundscapes/The vOICe are described at www.seeingwithsound.com
- Melnick, M.D., Tadin, D., and Huxlin, K.R.** (2016). Re-learning to see in cortical blindness. *Neuroscientist, 22*(2), 199–212.
- Meltzoff, A.N.** (1988). Imitation, objects, tools, and the rudiments of language in human ontogeny. *Human Evolution, 3*, 45–64.
- Meltzoff, A.N.** (1996). The human infant as imitative generalist: A 20-year progress report on infant imitation with implications for comparative psychology. In C.M. Heyes and B.G. Galef (Eds), *Social learning in animals: The roots of culture* (pp. 347–370). San Diego, CA: Academic Press.
- Melzack, R.** (1989). Phantom limbs, the self and the brain. *Canadian Psychology, 30*, 1–16.
- Melzack, R.** (1992). Phantom limbs. *Scientific American, 266*, 90–96.
- Menary, R.** (2010). Introduction to the special issue on 4E cognition. *Phenomenology and the Cognitive Sciences, 9*(4), 459–463.
- Meng, M., and Tong, F.** (2004). Binocular rivalry and perceptual filling-in of visual phantoms in human visual cortex. *Journal of Vision, 4*(8), article 63.
- Menon, G. Rahman, I., Menon, S. and Dutton, G.** (2003). Complex visual hallucinations in the visually impaired: The Charles Bonnet Syndrome. *Survey of Ophthalmology, 48*, 58–72.
- Mercier, C.** (1888). *The nervous system and the mind*. London: Macmillan.
- Merikle, P.M.** (2000). Subliminal perception. In A.E. Kazdin (Ed.), *Encyclopedia of psychology* (pp. 497–499). Vol. 7. New York: Oxford University Press.

Merikle, P. (2007). Preconscious processing. In M. Velmans and S. Schneider (Eds), *The Blackwell companion to consciousness* (pp. 512–524). Oxford: Blackwell.

Merikle, P.M., and Daneman, M. (1996). Memory for unconsciously perceived events: Evidence from anesthetized patients. *Consciousness and Cognition*, 5, 525–541.

Merikle, P.M., Smilek, D., and Eastwood, J. D. (2001). Perception without awareness: Perspectives from cognitive psychology. *Cognition*, 79(1–2), 115–134. Reprinted in S. Dehaene (Ed.) (2002), *The cognitive neuroscience of consciousness* (pp. 115–134). Cambridge, MA: MIT Press.

Merker, B. (2007). Consciousness without a cerebral cortex: A challenge for neuroscience and medicine. *Behavioral and Brain Sciences*, 30, 63–121 (incl. commentaries and author's response).

Merker, B. (2013). Cortical gamma oscillations: The functional key is activation, not cognition. *Neuroscience and Biobehavioral Reviews*, 37, 401–417.

Merleau-Ponty, M. (1945/2002). *Phenomenology of perception* [*Phénoménologie de la perception*]. Trans. C. Smith. London: Routledge.

Merleau-Ponty, M. (1942/1965). *The structure of behaviour* [*La structure du comportement*]. Trans. A. L. Fischer. London: Methuen.

Metzinger, T. (1995a). Faster than thought: Wholeness, homogeneity and temporal coding. In T. Metzinger (Ed.), *Conscious experience* (pp. 425–461). Thorverton, Devon: Imprint Academic.

Metzinger, T. (Ed.) (1995b). *Conscious experience*. Thorverton, Devon: Imprint Academic.

Metzinger, T. (Ed.) (2000). *Neural correlates of consciousness*. Cambridge, MA: MIT Press.

Metzinger, T. (Ed.) (2003). *Being no one: The self-model theory of subjectivity*. Cambridge, MA: MIT Press.

Metzinger, T. (2005). Out-of-body experiences as the origin of the concept of a 'soul'. *Mind and Matter*, 31, 57–84.

Metzinger, T. (2009). *The ego tunnel: The science of the mind and the myth of the self*. New York: Basic Books.

Metzner, R. (Ed.) (1999). *Ayahuasca: Human consciousness and the spirits of nature*. New York: Thunder's Mouth Press.

Metzner, R., and Adamson, S. (2001). Using MDMA in healing, psychotherapy, and spiritual practice. In J. Holland (Ed.), *Ecstasy: The complete guide: A comprehensive look at the risks and benefits of MDMA* (pp. 182–207). Rochester, VT: Park Street Press.

• R E F E R E N C E S

- Michaelson, J.** (2013). *Evolving dharma: Meditation, Buddhism, and the next generation of enlightenment*. Berkeley, CA: North Atlantic Books.
- Mikulas, W.L.** (2007). Buddhism & Western psychology: Fundamentals of integration. *Journal of Consciousness Studies*, 14(4), 4–49.
- Miles, J.B.** (2013). 'Irresponsible and a disservice': The integrity of social psychology turns on the free will dilemma. *British Journal of Social Psychology*, 52, 205–218.
- Miller, G.A.** (1962). *Psychology: the science of mental life*. New York: Harper & Row.
- Miller, S.M.** (2007). On the correlation/constitution distinction problem (and other hard problems) in the scientific study of consciousness. *Acta Neuropsychiatrica*, 19, 159–176.
- Milner, A.D.** (2008). Conscious and unconscious visual processing in the human brain. In L. Weiskrantz and M. Davies (Eds), *Frontiers of consciousness: Chichele Lectures* (pp. 169–214). Oxford: Oxford University Press.
- Milner, A. D.** (2012). Is visual processing in the dorsal stream accessible to consciousness? *Proceedings of the Royal Society of London B: Biological Sciences*, 279, 2289–2298.
- Milner, A.D., and Goodale, M.A.** (1995). *The visual brain in action*. Oxford: Oxford University Press.
- Milton, J.** (1999). Should ganzfeld research continue to be crucial in the search for a replicable psi effect? Part 1. Discussion paper and introduction to electronic-mail discussion. *Journal of Parapsychology*, 63, 309–333.
- Milton, J. and Wiseman, R.** (1999). Does psi exist? Lack of replication of an anomalous process of information transfer. *Psychological Bulletin*, 125, 387–391.
- Minsky, M.** (1986). *Society of mind*. New York: Simon and Schuster.
- Mithen, S.** (1996). *The prehistory of the mind: A search for the origins of art, religion and science*. London: Thames & Hudson.
- Mitchell, S.W.** (1871). Phantom limbs. *Lippincotts Magazine*, 8, 563–569.
- Mithoefer, M.C., Wagner, M.T., Mithoefer, A.T., Jerome, L., Martin, S.F., Yazar-Klosinski, B., Michel, Y., Brewerton, T.D., and Doblin, R.** (2013). Durability of improvement in post-traumatic stress disorder symptoms and absence of harmful effects or drug dependency after 3,4-methylenedioxymethamphetamine-assisted psychotherapy: A prospective long-term follow-up study. *Journal of Psychopharmacology*, 27(1), 28–39.
- Mitson, L., Ono, H., and Barbeito, R.** (1976). Three methods of measuring the location of the egocentre: Their reliability, comparative locations and intercorrelations. *Canadian Journal of Psychology*, 30, 1–8.

- Mohr, C., Binkofski, F., Erdmann, C., Büchel, C. and Helmchen, C.** (2009). The anterior cingulate cortex contains distinct areas dissociating external from self – administered painful stimulation: A parametric fMRI study. *Pain*, 114, 347–357.
- Mole, C.** (2010). *Attention is cognitive unison: An essay in philosophical psychology*. Oxford: Oxford University Press.
- Mole, C.** (2011). *Attention is cognitive unison: An essay in philosophical psychology*. New York: Oxford University Press.
- Moody, T.C.** (1994). Conversations with zombies. *Journal of Consciousness Studies*, 1(2), 196–200.
- Moody, T.C.** (1995). Why zombies won't stay dead. *Journal of Consciousness Studies*, 2(4), 365–372.
- Moorcroft, W.H.** (2013). *Understanding sleep and dreaming*. 2nd ed. Boston: Springer.
- Moore, A., and Malinowski, P.** (2009). Meditation, mindfulness and cognitive flexibility. *Consciousness and Cognition*, 18, 176–186.
- Moore, J.W., and Obhi, S.S.** (2012). Intentional binding and the sense of agency: A review. *Consciousness and Cognition*, 21(1), 546–561.
- Morewedge, C.K., and Norton, M.I.** (2009). When dreaming is believing: The (motivated) interpretation of dreams. *Journal of Personality and Social Psychology*, 96(2), 249–264.
- Morgan, C.J.A., and Curran, H.V., on behalf of the Independent Scientific Committee on Drugs (ISCD).** (2011). Ketamine use: A review. *Addiction*, 107, 27–38.
- Morgan, H.L., Turner, D.C., Corlett, P.R., Absalon, A.R., Adapa, R., Arana, F.S., Piggot, J., Gardner, J., Evertitt, J., Haggard, P., and Fletcher, P.C.** (2011). Exploring the impact of ketamine on the experience of illusory body ownership. *Biological Psychiatry*, 69(1), 35–41.
- Morland, A.B.** (1999). Conscious and veridical motion perception in a human hemianope. *Journal of Consciousness Studies*, 6(5), 43–53.
- Morris, R.L., Harary, S.B., Janis, J., Hartwell, J., and Roll, W.G.** (1978). Studies of communication during out-of-body experiences. *Journal of the American Society for Psychical Research*, 72, 1–22.
- Morse, M.** (1992). *Transformed by the light: The powerful effect of near-death experiences on people's lives*. New York: Villard.
- Morsella, E., Godwin, C.A., Jantz, T.K., Krieger, S.C., and Gazzaley, A.** (2016). Homing in consciousness in the nervous system: An

• R E F E R E N C E S

action-based synthesis. *Behavioral and Brain Sciences*, 39, e168–e199 (incl. commentaries and authors' response).

Most, S. B., Scholl, B. J., Clifford, E. R., and Simons, D. J. (2005). What you see is what you set: Sustained inattentional blindness and the capture of awareness. *Psychological Review*, 112, 217–242.

Moutoussis, K., and Zeki, S. (2002). The relationship between cortical activation and perception investigated with invisible stimuli. *Proceedings of the National Academy of Sciences, USA*, 99, 9527–9532.

Movshon, J. A. (2013). Three comments on Teller's 'bridge locus'. *Visual Neuroscience*, 30(0), 219–222.

Muetzelfeldt, L., Kamboj, S. K., Rees, H., Taylor, J., Morgan, C. J. A., and Curran, H. V. (2008). Journey through the K-hole: Phenomenological aspects of ketamine use. *Drug and Alcohol Dependence*, 95, 219–229.

Muldoon, S. J., and Carrington, H. (1929). *The projection of the astral body*. London: Rider & Co.

Müller, B. C. N., van Leeuwen, M. L. van Baaren, R. B., Bekkering, H., and Dijksterhuis, A. P. (2013). Empathy is a beautiful thing: Empathy predicts imitation only for attractive others. *Scandinavian Journal of Psychology*, 54, 401–406.

Müller, M. M., and Hübner, R. (2002). Can the spotlight of attention be shaped like a doughnut? Evidence from steady-state visual evoked potentials. *Psychological Science*, 13(2), 119–124.

Mullins, S., and Spence, S. A. (2003). Re-examining thought insertion. *The British Journal of Psychiatry*, 182(4), 293–298.

Murphy, S. T., and Zajonc, R. B. (1993). Affect, cognition, and awareness: Affective priming with optimal and suboptimal stimulus exposures. *Journal of Personality and Social Psychology*, 64, 723–739.

Murray, C. D. (Ed.) (2009). *Psychological scientific perspectives on out-of-body and near-death experiences*. New York: Buffalo.

Musso, F., Brinkmeyer, J., Ecker, D., London, M. K., Thieme, G., Warbrick, T., Wittsack, H.-J., Saleh, A., Greb, W., de Boer, P., and Winterer, G. (2011). Ketamine effects on brain function – Simultaneous fMRI/EEG during a visual oddball task. *NeuroImage*, 58, 508–525.

Myers, F. W. H. (1903). *Human personality and its survival of bodily death*. 2 vols. London: Longmans, Green, and Co.

Nagel, T. (1974). What is it like to be a bat? *Philosophical Review*, 83, 435–450. Reprinted with commentary in D. R. Hofstadter and D. C. Dennett (Eds) (1981), *The mind's I: Fantasies and reflections on self and soul* (pp. 391–414). London: Penguin. Also in Nagel, T. (1979), *Mortal questions* (pp. 165–180). Cambridge: Cambridge University Press. And in D. Chalmers (Ed.) (2002),

Philosophy of mind: Classical and contemporary readings (pp. 219–226). New York: Oxford University Press.

Nagel, T. (1979). *Mortal questions*. Cambridge: Cambridge University Press.

Nagel, T. (1986). *The view from nowhere*. New York: Oxford University Press.

Nelson, K.R., Mattingly, M., Lee, S.A., and Schmitt, F.A. (2006). Does the arousal system contribute to near death experience? *Neurology*, 66(7), 1003–1009.

Nietzsche, F. (1882/1974). *The gay science [Die fröhliche Wissenschaft]*. Trans. W. Kaufmann. New York: Random House.

Nishimoto, S., Vu, A.T., Naselaris, T., Benjamini, Y., Yu, B., and Gallant, J.L. (2011). Reconstructing visual experiences from brain activity evoked by natural movies. *Current Biology*, 21, 1641–1646.

Newberg, A., and D'Aquili, E. (2001). *Why God won't go away: Brain science and the biology of belief*. New York: Ballantine.

Nielsen, T.A. (2000). A review of mentation in REM and NREM sleep: 'Covert' REM sleep as a possible reconciliation of two opposing models. *Behavioral and Brain Sciences*, 23, 851–1121 (incl. commentaries and author's response).

Niesink, R.J.M., and van Laar, M.W. (2013). Does cannabidiol protect against adverse psychological effects of THC? *Frontiers in Psychology*, 4, article 130.

Nir, Y., and Tononi, G. (2010). Dreaming and the brain: From phenomenology to neurophysiology. *Trends in Cognitive Sciences*, 14(2), 88–100.

Noë, A. (Ed.). (2002). Is the visual world a grand illusion? Special issue, *Journal of Consciousness Studies*, 9(5–6). Also Thorverton, Devon: Imprint Academic.

Noë, A. (2005). *Action in perception*. Cambridge, MA: MIT Press.

Noë, A. (2009). *Out of our heads: Why you are not your brain, and other lessons from the biology of consciousness*. New York: Hill and Wang.

Noë, A., Pessoa, L., Thompson, E. (2000). Beyond the grand illusion: What change blindness really teaches us about vision. *Visual Cognition*, 7, 93–106.

Noë, A., and Thompson, E. (2004). Are there neural correlates of consciousness? *Journal of Consciousness Studies*, 11(1), 3–28.

Norman, L.J., Heywood, C.A., and Kentridge, R.W. (2013). Object-based attention without awareness. *Psychological Science*, 24(6), 836–843.

• R E F E R E N C E S

- Nolen-Hoeksema, S., Fredrickson, B. L., Loftus, G. R., and Lutz, C.** (2014). *Atkinson & Hilgard's introduction to psychology*. 16th edition. Andover: Cengage Learning.
- Nørretranders, T.** (1998). *The user illusion: Cutting consciousness down to size*. London: Penguin.
- Nowak, M., and Highfield, R.** (2011). *Supercooperators: Altruism, evolution, and why we need each other to succeed*. Edinburgh: Canongate.
- Nunn, C.** (Ed.) (2009). Defining consciousness. Special issue, *Journal of Consciousness Studies*, 16(5).
- Ogawa, K., Uema, T., Motohashi, N., Nishikawa, M., Takano, H., Hiroki, M., . . . and Yamada, Y.** (2003). Neural mechanism of propofol anesthesia in severe depression: A positron emission tomographic study. *Anesthesiology*, 98(5), 1101–1111.
- O'Hara, K. and Scutt, T.** (1996). There is no hard problem of consciousness. *Journal of Consciousness Studies*, 3(4), 290–302. Reprinted in J. Shear, ed. (1997), *Explaining consciousness – The 'hard problem'* (pp. 69–82). Cambridge, MA: MIT Press.
- Ólafsdóttir, H. F., Barry, C., Saleem, A. B., Hassabis, D., and Spiers, H. J.** (2015). Hippocampal place cells construct reward related sequences through unexplored space. *eLIFE*, 4, e06063.
- Olivares, F. A., Vargas, E., Fuentes, C., Martinez-Pernia, D., and Canales-Johnson, A.** (2015). Neurophenomenology revisited: Second-person methods for the study of human consciousness. *Frontiers in Psychology*, 6, 673.
- Olkowicz, S., Kocourek, M., Lučan, R. K., Porteš, M., Fitch, W. T., Herculano-Houzel, S., and Němec, P.** (2016). Birds have primate-like numbers of neurons in the forebrain. *Proceedings of the National Academy of Sciences, USA*, 113(26), 7255–7260.
- Olson, J. A., Amlani, A. A., Raz, A., and Rensink, R. A.** (2015). Influencing choice without awareness. *Consciousness and Cognition*, 37, 225–236.
- Olson, J. A., Landry, M., Appourchaux, K., and Raz, A.** (2016). Simulated thought insertion: Influencing the sense of agency using deception and magic. *Consciousness and Cognition*, 43, 11–26.
- O'Regan, J. K.** (1992). Solving the 'real' mysteries of visual perception: The world as an outside memory. *Canadian Journal of Psychology*, 46, 461–488.
- O'Regan, J. K.** (2011). *Why red doesn't sound like a bell: Understanding the feel of consciousness*. New York: Oxford University Press.
- O'Regan, J. K., and Noë, A.** (2001). A sensorimotor account of vision and visual consciousness. *Behavioral and Brain Sciences*, 24, 939–1031 (incl. commentaries and authors' response).

O'Regan, J.K., Rensink, R.A., and Clark, J.J. (1999). Change-blindness as a result of 'mudsplashes'. *Nature*, 398, 34.

Ornstein, R.E. (1986). *The psychology of consciousness*. 3rd ed. New York: Penguin.

Ornstein, R.E. (1991). *The evolution of consciousness: Of Darwin, Freud, and cranial fire: The origins of the way we think*. New York: Prentice Hall.

Otero-Millan, J., Macknik, S.L., Robbins, A., and Martinez-Conde, S. (2011). Stronger misdirection in curved than in straight motion. *Frontiers in Human Neuroscience*, 5, article 133.

Ott, U. (2001). The EEG and the depth of meditation. *Journal for Meditation and Meditation Research*, 1, 55–68.

Pace-Schott, E.F., and Hobson, J.A. (2007). Altered states of consciousness: Drug-induced states. In M. Velmans and S. Schneider (Eds), *The Blackwell companion to consciousness* (pp. 141–153). Oxford: Blackwell.

Pahnke, W. (1963). *Drugs and mysticism: An analysis of the relationship between psychedelic drugs and the mystical consciousness*. PhD thesis, Harvard University.

Pahnke, W. (1967). LSD and religious experience. In R. DeBold and R. Leaf (Eds), *LSD, man and society* (pp. 60–85). Middletown, CT: Wesleyan University Press.

Paladino, M.-P., Mazurega, M., Pavani, F., and Schubert, T.W. (2010). Synchronous multisensory stimulation blurs self-other boundaries. *Psychological Science*, 21(9), 1202–1207.

Palagi, E., Leone, A., Mancini, G., and Ferrari, P.F. (2009). Contagious yawning in gelada baboons as a possible expression of empathy. *Proceedings of the National Academy of Sciences, USA*, 106(46), 19262–19267.

Paley, W. (1802). *Natural theology: Or, evidences of the existence and attributes of the deity, collected from the appearances of nature*. London: Charles Knight. (Page numbers are from 15th ed., 1815.)

Palmer, J. (1978). The out-of-body experience: A psychological theory. *Parapsychology Review*, 9(5), 19–22.

Palmer, J. (2003). ESP in the ganzfeld: Analysis of a debate. *Journal of Consciousness Studies*, 10(6–7), 51–68.

Papineau, D. (2002). *Thinking about consciousness*. Oxford: Oxford University Press.

Papineau, D. (2003a). Theories of consciousness. In Q. Smith and A. Jokic (Eds), *Consciousness: New philosophical perspectives* (pp. 353–383). New York: Oxford University Press.

Papineau, D. (2003b). Confusions about consciousness. *Richmond Journal of Philosophy*, 5.

• R E F E R E N C E S

- Parfit, D.** (1984). *Reasons and persons*. Oxford: Oxford University Press.
- Parfit, D.** (1987). Divided minds and the nature of persons. In C. Blakemore and S. Greenfield (Eds), *Mindwaves* (pp. 19–26). Oxford: Blackwell. Also reprinted in S. Schneider (Ed.) (2009), *Science fiction and philosophy: From time travel to superintelligence* (pp. 91–98). Chichester: Wiley-Blackwell.
- Parker, J.D., and Blackmore, S.J.** (2002). Comparing the content of sleep paralysis and dreams reports. *Dreaming: Journal of the Association for the Study of Dreams*, 12, 45–59.
- Parnia, S., and Fenwick, P.** (2002). Near death experiences in cardiac arrest: Visions of a dying brain or visions of a new science of consciousness. *Resuscitation*, 52, 5–11.
- Parnia, S., Spearpoint, K., de Vos, G., Fenwick, P., Goldberg, D., Yang, J., . . . and Schoenfeld, E.R.** (2014). AWARE – AWAreness during REscuscitation – A prospective study. *Resuscitation*, 85(12), 1799–1805.
- Parnia, S., Waller, D.G., Yeates, R., and Fenwick, P.** (2001). A qualitative and quantitative study of the incidence, features and aetiology of near death experiences in cardiac arrest survivors. *Resuscitation*, 48, 149–156.
- Parra, A.** (2010). Out-of-body experiences and hallucinatory experiences: A psychological approach. *Imagination, Cognition and Personality*, 29(3), 211–223.
- Parris, B.A., Kuhn, G., Mizon, G.A., Benattayallah, A., and Hodgson, T.L.** (2009). Imaging the impossible: An fMRI study of impossible causal relationships in magic tricks. *NeuroImage*, 45(3), 1033–1039.
- Pashler, H.** (1998). *The psychology of attention*. Cambridge, MA: MIT Press.
- Pauen, M., Staudacher, A., and Walter, S.** (2006). Epiphenomenalism: Dead end or way out? *Journal of Consciousness Studies*, 13(1–2), 7–19.
- Paukner, A., Suomi, S. J., Visalberghi, E., and Ferrari, P.F.** (2009). Capuchin monkeys display affiliation toward humans who imitate them. *Science*, 325(5942), 880–883.
- Paulhus, D.L., and Carey, J.M.** (2011). The FAD-Plus: measuring lay beliefs regarding free will and related constructs. *Journal of Personality Assessment*, 93(1), 96–104.
- Paulignan, Y., MacKenzie, C., Marteniuk, R., and Jeannerod, M.** (1990). The coupling of arm and finger movements during prehension. *Experimental Brain Research*, 79, 431–435.
- Pearson, J., and Kosslyn, S.M.** (2015). The heterogeneity of mental representation: Ending the imagery debate. *Proceedings of the National Academy of Sciences*, 112(33), 10089–10092.
- Peirce, C.S., and Jastrow, J.** (1885). On small differences of sensation. *Memoirs of the National Academy of Sciences*, 3, 75–83.

- Penfield, W.** (1955). The role of the temporal cortex in certain psychical phenomena. *The Journal of Mental Science*, 101, 451–465.
- Penn, D. C., and Povinelli, D. J.** (2007). On the lack of evidence that non-human animals possess anything remotely resembling a ‘theory of mind’. *Philosophical Transactions of the Royal Society B*, 362, 731–744.
- Pennartz, C.** (2009). Identification and integration of sensory modalities: Neural basis and relation to consciousness. *Consciousness and Cognition*, 18, 718–739.
- Pennartz, C.** (2015). *The brain’s representational power: On consciousness and the integration of modalities*. Cambridge, MA: MIT Press.
- Penrose, R.** (1989). *The emperor’s new mind: Concerning computers, minds and the laws of physics*. Oxford: Oxford University Press.
- Penrose, R.** (1994a). *Shadows of the mind: A search for the missing science of consciousness*. Oxford: Oxford University Press.
- Penrose, R.** (1994b). Mechanisms, microtubules and the mind. *Journal of Consciousness Studies*, 1(2), 241–249.
- Perky, C. W.** (1910). An experimental study of imagination. *American Journal of Psychology*, 21, 422–452.
- Perna, A., Tosetti, M., Montanaro, D., and Morrone, M. C.** (2005). Neuronal mechanisms for illusory brightness perception in humans. *Neuron*, 47, 645–651.
- Perry, E. K.** (2002). Plants of the gods: Ethnic routes to altered consciousness. In E. Perry, H. Ashton and A. Young (Eds), *Neurochemistry of consciousness: Neurotransmitters in mind* (pp. 205–225). Amsterdam: John Benjamins.
- Persinger, M. A.** (1983). Religious and mystical experiences as artifacts of temporal lobe function: A general hypothesis. *Perceptual and Motor Skills*, 57, 1255–1262.
- Persinger, M. A.** (1999). *Neuropsychological bases of God beliefs*. Westport, CT: Praeger.
- Pessiglione, M., Schmidt, L., Draganski, B., Kalisch, R., Lau, H., Dolan, R. J., and Frith, C. D.** (2007). How the brain translates money into force: A neuroimaging study of subliminal motivation. *Science*, 316, 904–906.
- Pessoa, L., Thompson, E., and Noë, A.** (1998). Finding out about filling-in: A guide to perceptual completion for visual science and the philosophy of perception. *Behavioral and Brain Sciences*, 21, 723–802 (incl. commentaries and authors’ response).
- Petitmengin, C., and Lachaux, J.-P.** (2013). Microcognitive science: Bridging experiential and neuronal microdynamics. *Frontiers in Human Neuroscience*, 7, article 617.

• R E F E R E N C E S

Petitmengin, C., Remillieux, A., Cahour, B., and Carter-Thomas, S. (2013). A gap in Nisbett's and Wilson's findings? A first-person access to our cognitive processes. *Consciousness and Cognition*, 22(2), 654–669.

Petkova, V.I., and Ehrsson, H.H. (2008). If I were you: Perceptual illusion of body swapping. *PLOS ONE*, 3(12), e3832.

Phillips, F., Natter, M.B., and Egan, E.J.L. (2015). Magically deceptive biological motion – the French Drop Sleight. *Frontiers in Psychology*, 6, article 371.

Piccinini, G. (2010). How to improve on heterophenomenology: The self-measurement methodology of first-person data. *Journal of Consciousness Studies*, 17(3–4), 84–106.

Pickering, J. (Ed.) (1997). *The authority of experience: Essays on Buddhism and psychology*. Richmond, Surrey: Curzon.

Pickering, J., and Skinner, M. (Eds) (1990). *From sentience to symbols: Readings on consciousness*. London: Harvester Wheatsheaf.

Piet, J., and Hougaard, E. (2011). The effect of mindfulness-based cognitive therapy for prevention of relapse in recurrent major depressive disorder: A systematic review and meta-analysis. *Clinical Psychology Review*, 31, 1032–1040.

Pigliucci, M. (2013). What hard problem? *Philosophy Now*, 99, November/December. https://philosophynow.org/issues/99/What_Hard_Problem

Pinker, S. (1994). *The language instinct*. New York: Morrow.

Pinker, S. (1997). *How the mind works*. New York: W.W. Norton.

Pinker, S. (2002). *The blank slate: The modern denial of human nature*. New York: Viking.

Pinker, S. (2007). The brain: The mystery of consciousness. *Time*, 29 January 2007. <http://content.time.com/time/magazine/article/0,9171,1580394-1,00.html>

Pinker, S. (2016). The false allure of group selection. In D.M. Buss (Ed.), *The handbook of evolutionary psychology. Volume 2: Integrations* (pp. 867–880). 2nd ed. Hoboken, NJ: Wiley.

Pino, S., and Di Mauro, E. (2014). How to conciliate Popper with Cartesius. Comment on: 'Consciousness in the universe. A review of the "Orch OR" theory' by S. Hameroff and R. Penrose. *Physics of Life Reviews*, 11, 91–93.

Plante, S.G. (2017). LSD microdoses make people feel sharper, and scientists want to know how. *The Verge*, 24 April. www.theverge.com/2017/4/24/15403644/microdosing-ldd-acid-productivity-benefits-brain-studies

Plotnik, J.M., de Waal, F.B.M., and Reiss, D. (2006). Self-recognition in an Asian elephant. *Proceedings of the National Academy of Sciences, USA*, 103, 17053–17057.

- Poldrack, E. A.** (2011). Inferring mental states from neuroimaging data: From reverse inference to large-scale decoding. *Neuron*, 72(5), 692–697.
- Politzer, P.** (2008). *Neuroeconomics: A guide to the new science of making choices*. New York: Oxford University Press.
- Pomarol-Clotet, E., Honey, G. D., Murray, G. K., Corlett, P. R., Absalom, A. R., Lee, M., . . . and Fletcher, P. C.** (2006). Psychological effects of ketamine in healthy volunteers. *The British Journal of Psychiatry*, 189(2), 173–179.
- Popper, K. R., and Eccles, C.** (1977). *The self and its brain: An argument for interactionism*. New York: Springer.
- Povinelli, D. J.** (1998). Can animals empathize? Maybe not. *Scientific American*, 9(4), 67–76.
- Povinelli, D. J.** (2001). The self: Elevated in consciousness and extended in time. In C. Moore and K. Lemmon (Eds), *The self in time: Developmental perspectives* (pp. 75–95). Mahwah, NJ: Lawrence Erlbaum.
- Premack, D., and Woodruff, G.** (1978). Does the chimpanzee have a theory of mind? *Behavioral and Brain Sciences*, 1(4), 515–629 (incl. commentaries and authors' response).
- Prettyman, A.** (2012). Empty thoughts: An explanatory problem for higher-order theories of consciousness. *Consciousness Online*, 17 February. <https://consciousnessonline.wordpress.com/2012/02/17/empty-thoughts-an-explanatory-problem-for-higher-order-theories-of-consciousness/>
- Price, D. D., and Barrell, J. J.** (2012). *Inner experience and neuroscience: Merging both perspectives*. Cambridge, MA: MIT Press.
- Prince, M.** (1906). *The dissociation of a personality*. New York: Longmans, Green, and Co.
- Prinz, J. J.** (2003). Level-headed mysterianism and artificial experience. *Journal of Consciousness Studies*, 10, 111–132.
- Prinz, J.** (2007). The intermediate level theory of consciousness. In M. Velmans and S. Schneider (Eds), *The Blackwell companion to consciousness* (pp. 247–260). Oxford: Blackwell.
- Prinz, J. J.** (2012). *The conscious brain: How attention engenders experience*. New York: Oxford University Press.
- Prinzmetal, W., Long, V., and Leonhardt, J.** (2008). Involuntary attention and brightness contrast. *Perception & Psychophysics*, 70(7), 1139–1150.
- Prior, H., Schwarz, A., and Güntürkün, O.** (2008). Mirror-induced behavior in the magpie (*Pica pica*): Evidence of self-recognition. *PLoS Biology*, 6(8), e202.

• R E F E R E N C E S

- Pronin, E., Wegner, D. M., McCarthy, K., and Rodriguez, S.** (2006). Everyday magical powers: the role of apparent mental causation in the overestimation of personal influence. *Journal of Personality and Social Psychology*, *91*, 218–231.
- Pylyshyn, Z. W.** (1973). What the mind's eye tells the mind's brain: A critique of mental imagery. *Psychological Bulletin*, *80*, 1–25.
- Pylyshyn, Z.** (2003). *Seeing and visualizing: It's not what you think*. Cambridge, MA: MIT Press.
- Radin, D. I., and Radin.** (1997). *The conscious universe: The scientific truth of psychic phenomena* (pp. 138–142). San Francisco, CA: HarperEdge.
- Raffman, D.** (1995). On the persistence of phenomenology. In T. Metzinger (Ed.), *Conscious experience* (pp. 293–308). Thorverton, Devon: Imprint Academic.
- Raffone, A., and Pantani, M.** (2010). A global workspace model for phenomenal and access consciousness. *Consciousness and Cognition*, *19*(2), 580–596.
- Rahula, W.** (1959). *What the Buddha taught*. London: Gordon Fraser, and New York: Grove Press.
- Raichle, M. E., MacLeod, A. M., Snyder, A. Z., Powers, W. J., Gusnard, D. A., and Shulman, G. L.** (2001). A default mode of brain function. *Proceedings of the National Academy of Sciences, USA*, *98*(2), 676–682.
- Rakos, R. F., Laurene, K. R., Skala, S., and Slane, S.** (2008). Belief in free will: Measurement and conceptualization innovations. *Behavior and Social Issues*, *17*, 20–39.
- Ramachandran, V. S., and Blakeslee, S.** (1998). *Phantoms in the brain*. London: Fourth Estate.
- Ramachandran, V. S., and Gregory, R. L.** (1991). Perceptual filling in of artificially induced scotomas in human vision. *Nature*, *350*, 699–702.
- Ramachandran, V. S., and Hirstein, W.** (1997). Three laws of qualia: What neurology tells us about the biological functions of consciousness. *Journal of Consciousness Studies*, *4*(5–6), 429–457.
- Ramachandran, V. S., and Hubbard, E. M.** (2001). Synaesthesia – A window into perception, thought and language. *Journal of Consciousness Studies*, *8*, 3–34.
- Rangarajan, V., Hermes, D., Foster, B. L., Weiner, K. S., Jacques, C., Grill-Spector, K., and Parvizi, J.** (2014). Electrical stimulation of the left and right human fusiform gyrus causes different effects in conscious face perception. *The Journal of Neuroscience*, *34*(38), 12828–12836.
- Rangel, A., Camerer, C., and Montague, R.** (2008). A framework for studying the neurobiology of value-based decision-making. *Nature Reviews Neuroscience*, *9*, 545–556.

- Ransom, M., Fazelpour, S., and Mole, C.** (2017). Attention in the predictive mind. *Consciousness and Cognition*, 47, 99–112.
- Recanzone, G. H.** (2009). Interactions of auditory and visual stimuli in space and time. *Hearing Research*, 258(1–2), 89–99.
- Reid, T.** (1785). *Essays on the intellectual powers of man*. Edinburgh: John Bell.
- Reinerth, M. S., and Thon, J.-N.** (2016). *Subjectivity across media: Interdisciplinary and transmedial perspectives*. New York: Routledge.
- Reiss, D.** (1998). Cognition and communication in dolphins: A question of consciousness. In S. R. Hameroff, A. W. Kaszniak, and A. C. Scott (Eds), *Toward a science of consciousness II: The Second Tucson Discussions and Debates* (pp. 551–560). Cambridge, MA: MIT Press.
- Reiss, D., and Marino, L.** (2001). Mirror self – recognition in the bottlenose dolphin: a case of cognitive convergence. *Proceedings of the National Academy of Sciences, USA*, 98, 5937–5942.
- Rensink, R.** (2000). The dynamic representation of scenes. *Visual Cognition*, 7, 17–42.
- Rensink, R. A., and Kuhn, G.** (2015). A framework for using magic to study the mind. *Frontiers in Psychology*, 5, article 1508.
- Rensink, R. A., O'Regan, J. K., and Clark, J. J.** (1997). To see or not to see: The need for attention to perceive changes in scenes. *Psychological Science*, 8(5), 368–373.
- Revonsuo, A.** (1999). Binding and the phenomenal unity of consciousness. *Consciousness and Cognition*, 8, 173–185.
- Revonsuo, A.** (2000). The reinterpretation of dreams: An evolutionary hypothesis of the function of dreaming. *Behavioral and Brain Sciences*, 23(6), 877–1121 (incl. commentaries and author's response).
- Revonsuo, A.** (2009). Binding problem. In T. Bayne, A. Cleeremans, and P. Wilken (Eds), *The Oxford companion to consciousness* (pp. 101–105). Oxford: Oxford University Press.
- Revonsuo, A., Kallio, S., and Sikka, P.** (2009). What is an altered state of consciousness? *Philosophical Psychology*, 22(2), 187–204.
- Revonsuo, A., and Tarkko, K.** (2002). Binding in dreams: The bizarreness of dream images and the unity of consciousness. *Journal of Consciousness Studies*, 9(7), 3–24.
- Richerson, P. J., and Boyd, R.** (2005). *Not by genes alone: How culture transformed human evolution*. Chicago: University of Chicago Press.
- Ridley, Mark** (1996). *Evolution*. 2nd ed. Oxford: Blackwell Science.
- Ridley, Matt** (1996). *The origins of virtue*. London: Viking.

• R E F E R E N C E S

Rilke, R. M. (1902). The panther [Der Panther]. Full text available at <http://www.picture-poems.com/rilke/panther.html>

Rilke, R. M. (1980). Selbstzeugnisse. In U. Fülleborn and M. Engel (Eds), *Materialien zu Rainer Maria Rilkes 'Duineser Elegien'*. Frankfurt am Main: Suhrkamp.

Ring, K. (1980). *Life at death: A scientific investigation of the near-death experience*. New York: Coward, McCann and Geoghegan.

Rizzolatti, G., and Craighero, L. (2010). Premotor theory of attention. *Scholarpedia*, 5(1), 6311. www.scholarpedia.org/article/Premotor_theory_of_attention

Rizzolatti, G., Riggio, L., and Sheliga, B.M. (1994). Space and selective attention. In C. Umiltà and M. Moscovitch (Eds), *Attention and performance XV: Conscious and nonconscious information processing* (pp. 231–265). Cambridge, MA: MIT Press.

Robbins, P., and Jack, A.I. (2006). The phenomenal stance. *Philosophical Studies*, 127, 59–85.

Rosch, E. (1997). Transformation of the wolf man. In J. Pickering (Ed.), *The authority of experience: Essays on Buddhism and psychology* (pp. 6–27). Richmond, Surrey: Curzon.

Rose, D. (2006). *Consciousness: Philosophical, psychological and neural theories*. Oxford: Oxford University Press.

Rose, H., and Rose, S. (2000). *Alas, poor Darwin: Arguments against evolutionary psychology*. London: Jonathan Cape.

Rosenberg, M.D., Finn, E.S., Scheinost, D., Constable, R.T., and Chun, M.M. (2017). Characterizing attention with predictive network models. *Trends in Cognitive Sciences*, 21(4), 290–302.

Rosenthal, D.M. (1995). Multiple drafts and the facts of the matter. In T. Metzinger (Ed.), *Conscious experience* (pp. 359–372). Thorverton, Devon: Imprint Academic.

Rosenthal, D.M. (2008). Consciousness and its function. *Neuropsychologia*, 46, 829–840.

Ross, J., Yilmaz, M., Dale, R., Cassidy, R., Yildirim, I., Zeedyk, M.S. (2017). Cultural differences in self-recognition: The early development of autonomous and related selves? *Developmental Science*, 20(3), e12387.

Ross, S., Bossis, A., Guss, J., Agin-Liebes, G., Malone, T., Cohen, B., Mennenga, S.E., Belser, A., Kalliontzi, K., Babb, J., Su, Z., Corby, P., and Schmidt, B.L. (2016). Rapid and sustained symptom reduction following psilocybin treatment for anxiety and depression in patients with life-threatening cancer: A randomized controlled trial. *Journal of Psychopharmacology*, 30(1), 1165–1180.

- Rothkirch, M., and Hesselmann, G.** (2017). What we talk about when we talk about unconscious processes – a plea for best practices. *Frontiers in Psychology*, 8, article 835.
- Rousseau, J. -J.** (1782–1789). *Confessions [Les Confessions]*. Paris: Cazin. Full text available at <https://www.gutenberg.org/files/3913/3913-h/3913-h.htm> and <https://books.google.co.uk/books?id=TyVbAAAAQAAJ>; also <http://www.lettres.org/confessions/confessions.htm> (original French)
- Rowlandson, W.** (2012). Nourished by dreams, visions, and William James: the radical philosophies of Borges and Terence McKenna. *Paranthropology*, 3(1), 46–60.
- Ruff, C.** (2011). A systems-neuroscience view of attention. In C. Mole, D. Smithies and W. Wu (Eds), *Attention: Philosophical and psychological essays* (pp. 1–23). Oxford: Oxford University Press.
- Ruggieri, V., and Alfieri, G.** (1992). The eyes in imagery and perceptual processes: First remarks. *Perceptual and Motor Skills*, 75, 287–290.
- Russell, B.** (1914). On the nature of acquaintance. II. Neutral monism. *The Monist*, 24(2), 161–187.
- Ryle, G.** (1949/2009). *The concept of mind*. New York: Routledge.
- Sacks, O.** (1985). *The man who mistook his wife for a hat, and other clinical tales*. London: Duckworth.
- Sagan, C.** (1971). Mr. X. In L. Grinspoon, *Marihuana reconsidered* (Ch. 4). Boston, MA: Harvard University Press.
- Sagan, C.** (2006). *The varieties of scientific experience: A personal view of the search for God*. New York: Penguin.
- Salva, O.R., Rugani, R., Cavazzana, A., Regolin, L., and Vallortigara, G.** (2013). Perception of the Ebbinghaus illusion in four-day-old chicks (*Gallus gallus*). *Animal Cognition*, 16(6), 895–906.
- Sand, N.** (2014, 28 April). Moving into the sacred world of DMT. *Psychedelic Frontier*.
- Sargent, C.** (1987). Sceptical fairytales from Bristol. *Journal of the Society for Psychical Research*, 54, 208–218.
- Sarkissian, H., Chatterjee, A., de Brigard, F., Knobe, J., Nichols, S., and Sirker, S.** (2010). Is belief in free will a cultural universal? *Mind & Language*, 25(3), 346–358.
- Sartre, J.-P.** (1940/2004). *The imaginary: A phenomenological psychology of the imagination [L'Imaginaire: Psychologie phénoménologique de l'imagination]*. Trans. J. Webber. London: Routledge.

• R E F E R E N C E S

- Sato, J.R., Kozasa, E. H., Russell, T. A., Radvany, J., Mello, L. E., Lacerda, S. S., and Amaro, E., Jr.** (2012). Brain imaging analysis can identify participants under regular mental training. *PLOS ONE*, 7(7), 1–6.
- Saunders, G.** (2014). *Acts of consciousness: A social psychology standpoint*. Cambridge: Cambridge University Press.
- Saunders, N.** (1993). *E for ecstasy*. London: Nicholas Saunders.
- Saunders, N., Saunders, A., and Pauli, M.** (2000). *In search of the ultimate high: Spiritual experiences through psychoactives*. London: Rider.
- Schacter, D.L., Addis, D. R., Hassabis, D., Martin, V.C., Spreng, N., and Szpunar, K.K.** (2012). The future of memory: Remembering, imagining, and the brain. *Neuron*, 21, 76(4).
- Schechtman, M.** (2011). The narrative self. In S. Gallagher (Ed.), *The Oxford handbook of the self* (pp. 394–416). New York: Oxford University Press.
- Scheidegger, M., Walter, M., Lehmann, M., Metzger, C., Grimm, S., Boeker, H., Boesiger, P., Henning, A., and Seifritz, E.** (2012). Ketamine decreases resting state functional network connectivity in healthy subjects: Implications for antidepressant drug action. *PLOS ONE*, 7(9), e44799.
- Schiff, N.D.** (2007). Global disorders of consciousness. In M. Velmans and S. Schneider (Eds), *The Blackwell companion to consciousness* (pp. 589–604). Oxford: Blackwell.
- Schilbach, L., Timmermans, B., Reddy, V., Costall, A., Bente, G., Schlicht, T., and Vogeley, K.** (2013). Toward a second-person neuroscience? *Behavioral and Brain Sciences*, 36(4), 393–462 (incl. commentaries and authors' response).
- Schneider, W.** (2009). Automaticity and consciousness. In W.P. Banks (Ed.), *Encyclopedia of consciousness* (pp. 83–92). New York: Academic.
- Schofield, T.P., Creswell, J.D., and Denson, T.F.** (2015). Brief mindfulness induction reduces inattention blindness. *Consciousness and Cognition*, 37, 63–70.
- Scholl, B.J., Noles, N.S., Pasheva, V., and Sussman, R.** (2003). Talking on a cellular telephone dramatically increases 'sustained inattention blindness' [Vision Sciences Society Annual Meeting Abstract]. *Journal of Vision*, 3(9), 156.
- Schultz, W.** (1999). The primate basal ganglia and the voluntary control of behaviour. *Journal of Consciousness Studies*, 6(8–9), 31–45. Reprinted in B. Libet, A. Freeman, and K. Sutherland (Eds), *The volitional brain: Towards a neuroscience of free will* (pp. 31–45). Thorverton, Devon: Imprint Academic.
- Schultze-Kraft, M., Birman, D., Rusconi, M., Allefeld, C., Görden, K., Dähne, S., Blankertz, B., and Haynes, J.-D.** (2015).

The point of no return in vetoing self-initiated movements. *Proceedings of the National Academy of Sciences, USA*, 113(4), 1080–1085.

Schurger, A., Sitt, J.D., and Dehaene, S. (2012). An accumulator model for spontaneous neural activity prior to self-initiated movement. *Proceedings of the National Academy of Sciences, USA*, 109(42), E2904–E2913.

Schwitzgebel, E. (2008). The unreliability of naïve introspection. *Philosophical Review*, 117(2), 245–273.

Seager, W. (2016). *Theories of consciousness: An introduction and assessment*. 2nd ed. London: Routledge.

Searle, J. (1980). Minds, brains, and programs. *Behavioral and Brain Sciences*, 3(3), 417–457 (incl. commentaries and author's response).

Searle, J. R. (1984). *Minds, brains and science*. Cambridge, MA: Harvard University Press.

Searle, J. (1992). *The rediscovery of the mind*. Cambridge, MA: MIT Press.

Searle, J. (1997). *The mystery of consciousness*. New York: New York Review of Books.

Searle, J. (1998). How to study consciousness scientifically. In S. R. Hameroff, A. W. Kazniak, and A. C. Scott (Eds), *Toward a science of consciousness II: The second Tucson discussions and debates* (pp. 15–30). Cambridge, MA: MIT Press.

Searle, J. (1999). I married a computer. Review of R. Kurzweil, *The age of spiritual machines: When computers exceed human intelligence*. *New York Review of Books*, 8 April.

Searle, J. R. (2002). Why I am not a property dualist. *Journal of Consciousness Studies*, 9(12), 57–64.

Segal, S. R. (Ed.) (2003). *Encountering Buddhism: Western psychology and Buddhist teachings*. New York: State University of New York Press.

Seth, A. (2007). Models of consciousness. *Scholarpedia*, 2(1), 1328. www.scholarpedia.org/article/Models_of_consciousness

Seth, A., Baars, B. J., and Edelman, D. B. (2005). Criteria for consciousness in humans and other mammals. *Consciousness and Cognition*, 14, 119–139.

Seth, A. K., Suzuki, K., and Critchley, H. D. (2012). An interoceptive predictive coding model of conscious presence. *Frontiers in Psychology*, 2, 395.

Sewell, A. (1877). *Black Beauty: The autobiography of a horse*. London: Jarrold and Sons. Full text available at <http://www.gutenberg.org/files/271/271-h/271-h.htm>

• R E F E R E N C E S

- Shadlen, M.N., and Movshon, J.A.** (1999). Synchrony unbound: A critical evaluation of the temporal binding hypothesis. *Neuron*, 24, 67–77.
- Shakespeare, W.** (1606). *King Lear*. (First quarto edition London 1608.) Full text available at <http://shakespeare.mit.edu/lear/full.html> and <https://books.google.co.uk/books?id=WvgqAAAAMAAJ>
- Shalom, D.E., de Sousa Serro, M.G., Giaconia, M., Martinez, L.M., Rieznik, A., and Sigman, M.** (2013). Choosing in freedom or forced to choose? Introspective blindness to psychological forcing in stage-magic. *PLOS ONE*, 8(3), e58254.
- Shanahan, M.** (2006). Towards a computational account of reflexive consciousness. *Proceedings of AISB'06: Adaptation in Artificial and Biological Systems*, 1, 165–170.
- Shanon, B.** (2002). *The antipodes of the mind: Charting the phenomenology of the ayahuasca experience*. Oxford: Oxford University Press.
- Shariff, A.F., Greene, J.D., Karremans, J.C., Luguri, J.B., Clark, C.J., Schooler, J.W., Baumeister, R. F., and Vohs, K.D.** (2014). Free will and punishment: A mechanistic view of human nature reduces retribution. *Psychological Science*, 25(8), 1563–1570.
- Sharpless, B.A., and Barber, J.P.** (2011). Lifetime prevalence rates of sleep paralysis: A systematic review. *Sleep Medicine Reviews*, 15(5), 311–315.
- Shear, J.** (Ed.) (1997). *Explaining consciousness – The 'hard problem'*. Cambridge, MA: MIT Press.
- Sheinberg, D.L., and Logothetis, N.K.** (1997). The role of temporal cortical areas in perceptual organization. *Proceedings of the National Academy of Sciences USA*, 94, 3408–3413.
- Sheldrake, R.** (2005). The sense of being stared at, Parts 1 and 2. *Journal of Consciousness Studies*, 12(6), 10–31, 32–49 (with open peer commentary).
- Sheng-Yen, Crook, J., Child, S., Kalin, M., and Andricevic, Z.** (2002). *Chan comes West*. Elmhurst, New York: Dharma Drum.
- Shepard, R.N., and Metzler, J.** (1971). Mental rotation of three-dimensional objects. *Science*, 171, 701–703.
- Sherry, D. F., and Galef, B. G.** (1984). Cultural transmission without imitation: Milk bottle opening by birds. *Animal Behavior*, 32, 937–938.
- Shifman, L.** (2013). Memes in a digital world: Reconciling with a conceptual troublemaker. *Journal of Computer-Mediated Communication*, 18(3), 362–377.
- Shulgin, A., and Shulgin, A.** (1991). *PiHKAL* (Phenethylamines I Have Known and Loved). Berkeley, CA: Transform Press.
- Shushruth, S.** (2013). Exploring the neural basis of consciousness through anesthesia. *The Journal of Neuroscience*, 33(5), 1757–1758.

Sidgwick, H., Sidgwick, E. M., and Johnson, A. (1894). Report on the census of hallucinations. *Proceedings of the Society for Psychical Research*, 10, 25–422.

Sidis, B. (1898). *The psychology of suggestion: A research into the subconscious nature of man and society*. New York: Appleton.

Siegel, R. K. (1977). Hallucinations. *Scientific American*, 237, 132–140.

Siegel, R. K. (1992). *Fire in the brain: Clinical tales of hallucination*. New York: Penguin.

Siegel, S. (2009). Contents of consciousness. In T. Bayne, A. Cleeremans, and P. Wilken (Eds), *The Oxford companion to consciousness* (pp. 189–192). Oxford: Oxford University Press.

Siegel, R. K., and Jarvik, M. E. (1975). Drug-induced hallucinations in animals and man. In R. K. Siegel and L. J. West (Eds), *Hallucinations: Behavior, experience, and theory* (pp. 81–161). New York: Wiley.

Simons, D. J. (2000). Current approaches to change blindness. *Visual Cognition*, 7, 1–15.

Simons, D. J., and Ambinder, M. S. (2005). Change blindness: Theory and consequences. *Current Directions in Psychological Science*, 14, 44–48.

Simons, D. J., and Chabris, C. F. (1999). Gorillas in our midst: Sustained inattention blindness for dynamic events. *Perception*, 28, 1059–1074.

Simons, D. J., Franconeri, S. L., and Reimer, R. L. (2000). Change blindness in the absence of a visual disruption. *Perception*, 29, 1143–1154.

Simons, D. J., and Levin, D. T. (1997). Change blindness. *Trends in Cognitive Sciences*, 1, 261–267.

Simons, D. J., and Levin, D. T. (1998). Failure to detect changes to people during a real-world interaction. *Psychonomic Bulletin and Review*, 5, 644–649.

Simons, D. J., and Rensink, R. A. (2005). Change blindness: Past, present, and future. *Trends in Cognitive Sciences*, 9(1), 16–20.

Singer, W. (2000). Phenomenal awareness and consciousness from a neurobiological perspective. In T. Metzinger (Ed.), *Neural correlates of consciousness: Empirical and conceptual questions* (pp. 121–137). Cambridge, MA: MIT Press.

Singer, W. (2007). Large-scale temporal coordination of cortical activity as a prerequisite for conscious experience. In M. Velmans and S. Schneider (Eds), *The Blackwell companion to consciousness* (pp. 605–615). Oxford: Blackwell.

Skinner, B. F. (1948). *Walden two*. New York: Macmillan.

Slade, P. D., and Bentall, R. P. (1988). *Sensory deception: A scientific analysis of hallucination*. Baltimore: Johns Hopkins University Press.

• R E F E R E N C E S

- Sligte, I. G., Scholte, H. S., and Lamme, V.A.F.** (2009). V4 activity predicts the strength of visual short-term memory representations. *Journal of Neuroscience*, 29(23), 7432–7438.
- Slogar, S.-M.** (2011). Dissociative identity disorder: Overview of research. *Inquiries Journal/Student Pulse*, 3(5). www.inquiriesjournal.com/a?id=525
- Sloman, A.** (2014). *Jane Austen's concept of information (Not Claude Shannon's)*. <https://www.cs.bham.ac.uk/research/projects/cogaff/misc/austen-info.html>
- Sloman, A., and Chrisley, R.** (2003). Virtual machines and consciousness. *Journal of Consciousness Studies*, 10, 133–172.
- Smit, R. H.** (2008). Corroboration of the dentures anecdote involving veridical perception in a near-death experience. *Journal of Near-Death Studies*, 27(1), 47–61.
- Smith, D. T., and Schenk, T.** (2012). The premotor theory of attention: Time to move on? *Neuropsychologia*, 50, 1104–1114.
- Smithies, D.** (2011). Attention is rational-access consciousness. In C. Mole, D. Smithies, and W. Wu (Eds), *Attention: Philosophical and psychological essays* (pp. 247–273). New York: Oxford University Press.
- Solms, M.** (2000). Dreaming and REM sleep are controlled by different brain mechanisms. *Behavioral and Brain Sciences*, 23(6), 843–1121 (incl. commentaries and author's response).
- Soon, C. S., Brass, M., Heinze, H.-J., and Haynes, J.-D.** (2008). Unconscious determinants of free decisions in the human brain. *Nature Neuroscience*, 11(5), 543–545.
- Sovrano, V.A., Albertazzi, L., and Salva, O.R.** (2014). The Ebbinghaus illusion in a fish (*Xenotoca eiseni*). *Animal Cognition*, 18(2), 533–542.
- Spanos, N.P.** (1991). A sociocognitive approach to hypnosis. In S. J. Lynn and J. W. Rhue (Eds), *Theories of hypnosis: Current models and perspectives* (pp. 324–362). New York: Guilford Press.
- Spence, S.A., and Frith, C.D.** (1999). Towards a functional anatomy of volition. *Journal of Consciousness Studies*, 6(8–9), 11–29. Reprinted in B. Libet, A. Freeman, and K. Sutherland (Eds), *The volitional Brain: Towards a neuroscience of free will* (pp. 11–29). Thorverton, Devon: Imprint Academic.
- Spring, M., and Carrasco, M.** (2015). Acting without seeing: Eye movements reveal visual processing without awareness. *Trends in Neurosciences*, 38(4), 247–258.
- Sperling, G.** (1960). The information available in brief visual presentations. *Psychological Monographs: General and Applied*, 74(11), 1–29.
- Sperry, R. W.** (1968). Hemisphere disconnection and unity in conscious awareness. *American Psychologist*, 23, 723–733.

- Speth, J., Speth, C., Kaelen, M., Schloerscheidt, A. M., Feilding, A., Nutt, D. J., and Carhart-Harris, R. L.** (2016). Decreased mental time travel to the past correlates with default-mode network disintegration under lysergic acid diethylamide. *Journal of Psychopharmacology*, 30(4), 344–353.
- Standage, T.** (2002). *The mechanical Turk: The true story of the chess-playing machine that fooled the world*. London: Penguin.
- Staniloiu, A., and Markowitsch, H. J.** (2014). Dissociative amnesia. *Lancet Psychiatry*, 1, 226–241.
- Stapp, H.** (2007). Quantum mechanical theories. In M. Velmans and S. Schneider (Eds), *The Blackwell companion to consciousness* (pp. 300–312). Oxford: Blackwell.
- Stapp, H. P.** (2011). *Mindful universe: Quantum mechanics and the participating observer*. 2nd ed. Berlin: Springer.
- Starmans, C., and Bloom, P.** (2012). Windows to the soul: Children and adults see the eyes as the location of the self. *Cognition*, 123(2), 313–318.
- Stazicker, J.** (2011). Attention, visual consciousness and indeterminacy. *Mind & Language*, 26(2), 156–184.
- Steels, L.** (2000). Language as a complex adaptive system. Lecture notes in computer science. In M. Schoenauer, K. Deb, G. Rudolph, X. Yao, E. Lutton, J. J. Merelo, and H.-P. Schwefel (Eds), *Parallel Problem Solving from Nature – PPSN-VI* (pp. 17–26). Berlin: Springer.
- Steels, L.** (2003). Language re-entrance and the ‘inner voice’. *Journal of Consciousness Studies*, 10(4–5), 173–185.
- Stein, B. E., Wallace, M. T., and Stanford, T. R.** (2001). Brain mechanisms for synthesizing information from different sensory modalities. In E. B. Goldstein and M. A. Malden (Eds), *Blackwell handbook of perception* (pp. 709–736). Oxford: Blackwell.
- Sterelny, K.** (2001). *Dawkins vs. Gould: Survival of the fittest*. Cambridge: Icon.
- Stevens, J.** (1987). *Storming heaven: LSD and the American dream*. New York: Atlantic Monthly Press.
- Stevens, R.** (2000). Phenomenological approaches to the study of conscious awareness. In M. Velmans (Ed.), *Investigating phenomenal consciousness* (pp. 99–120). Amsterdam: John Benjamins.
- Stevenson, R. L.** (1886). *The strange case of Dr Jekyll and Mr Hyde*. London: Longmans, Green, & Co.
- Stirling, J., and McCoy, L.** (2010). Quantifying the psychological effects of ketamine: From euphoria to the K-hole. *Substance Use & Misuse*, 45(14), 2428–2443.

• R E F E R E N C E S

- Strawson, G.** (1997). The self. *Journal of Consciousness Studies*, 4(5–6), 405–428. Reprinted in S. Gallagher and J. Shear (Eds) (1999), *Models of the self* (pp. 1–24). Exeter: Imprint Academic.
- Strawson, G.** (2006). Panpsychism?: Reply to commentators with a celebration of Descartes. *Journal of Consciousness Studies*, 13(10–11), 184–280.
- Strawson, G.** (2008). Realistic monism: Why physicalism entails panpsychism. In G. Strawson, *Real materialism and other essays* (pp. 53–74). Oxford: Clarendon Press.
- Strawson, G.** (2011). The minimal subject. In S. Gallagher (Ed.), *The Oxford handbook of the self* (pp. 253–278). New York: Oxford University Press.
- Stuart, S.A.J.** (2007). Machine consciousness: Cognitive and kinaesthetic imagination. *Journal of Consciousness Studies*, 14, 141–153.
- Stuart, S.A.J.** (2011). Enkinaesthesia: The fundamental challenge for machine consciousness. *International Journal of Machine Consciousness*, 3(1), 145–162.
- Studerus, E., Kometer, M., Hasler, F., and Vollenweider, F.X.** (2011). Acute, subacute and long-term subjective effects of psilocybin in healthy humans: A pooled analysis of experimental studies. *Journal of Psychopharmacology*, 25, 1434–1452.
- Stumbrys, T., and Erlacher, D.** (2016). Applications of lucid dreams and their effects on the mood upon awakening. *International Journal of Dream Research*, 9(2), 146–150.
- Stumbrys, T., Erlacher, D., and Schredl, M.** (2016). Effectiveness of motor practice in lucid dreams: A comparison with physical and mental practice. *Journal of Sports Sciences*, 34(1), 27–34.
- Suddendorf, T., and Butler, D.L.** (2013). The nature of visual self-recognition. *Trends in Cognitive Sciences*, 17(3), 121–127.
- Sully, J.** (1892). *The human mind: A text-book of psychology*. 2 vols. London: Longmans, Green & Co.
- Sur, M., and Leamey, C.** (2001). Development and plasticity of cortical areas and networks. *Nature Reviews Neuroscience*, 2(4), 251–262.
- Sutherland, K.** (Ed.) (1995). Zombie earth: Editorial introduction to a symposium on Todd Moody's 'conversations with zombies'. *Journal of Consciousness Studies*, 2(4), 312–372.
- Szegedy, C., Zaremba, W., Sutskever, I., Bruna, J., Erhan, D., Goodfellow, I., and Fergus, R.** (2014). Intriguing properties of neural networks. arXiv:1312.6199v4, 19 February.
- Tagliazucchi, E., Roseman, L., Kaelen, M., Orban, C., Muthukumaraswamy, S.D., Murphy, K., . . . and Bullmore, E.** (2016).

Increased global functional connectivity correlates with LSD-induced ego dissolution. *Current Biology*, 26(8), 1043–1050.

Tajadura-Jiménez, A., Longo, M. R., Coleman, R., and Tsakiris, M. (2012). The person in the mirror: Using the enfacement illusion to investigate the experiential structure of self-identification. *Consciousness and Cognition*, 21, 1725–1738.

Tallon-Baudry, C. (2003). Oscillatory synchrony as a signature for the unity of visual experience in humans. In A. Cleeremans (Ed.), *The unity of consciousness: Binding, integration and dissociation* (pp. 153–167). New York: Oxford University Press.

Tallon-Baudry, C., and Bertrand, O. (1999). Oscillatory gamma activity in humans and its role in object representation. *Trends in Cognitive Science*, 3(4), 151–162.

Targ, R. (2004). *Limitless mind: A guide to remote viewing and transformation of consciousness*. Novato, CA: New World Library.

Targ, R., and Puthoff, H. (1977). *Mind-reach: Scientists look at psychic abilities*. New York: Delacorte.

Tart, C. T. (1968). A psychophysiological study of out-of-body experiences in a selected subject. *Journal of the American Society for Psychical Research*, 62, 3–27.

Tart, C. T. (1971). *On being stoned: A psychological study of marijuana intoxication*. Palo Alto, CA: Science and Behavior Books.

Tart, C. T. (1972a). States of consciousness and state-specific sciences. *Science*, 176, 1203–1210.

Tart, C. T. (1972b). Concerning the scientific study of the human aura. *Journal of the Society for Psychical Research*, 46(751), 1–21.

Tart, C. T. (1975). *States of consciousness*. New York: Dutton & Co.

Tart, C. T. (2015). State-specific sciences: Altered state origin of the proposal. Charles C. Tart: Home page and blog, 7 September. <http://blog.paradigm-sys.com/state-specific-sciences-altered-state-origin-of-the-proposal/>

Tatler, B. W., and Land, M. F. (2011). Vision and the representation of the surroundings in spatial memory. *Philosophical Transactions of the Royal Society B*, 366, 596–610.

Taylor, J. L., and McCloskey, D. I. (1990). Triggering of preprogrammed movements as reactions to masked stimuli. *Journal of Neurophysiology*, 63, 439–446.

Taylor, R. (Ed.) (1843). *Scientific memoirs, Selected from the transactions of foreign academies and learned societies and from foreign journals*. London: Richard and John E. Taylor.

• R E F E R E N C E S

- Taylor, S.** (2012). Transformation through suffering: A study of individuals who have experienced positive psychological transformation following periods of intense turmoil. *Journal of Humanistic Psychology*, 52(1), 30–52.
- Taylor, S.** (2017). *The leap: The psychology of spiritual awakening*. Eckhart Tolle edition. Carlsbad, CA: Hay House.
- Teller, D. Y., and Pugh Jr., E. N.** (1983). Linking propositions in color vision. In J. D. Mollon and L. T. Sharpe (Eds), *Colour vision: Physiology and psychophysics* (pp. 577–589). London: Academic.
- Terrace, H.** (1987). Thoughts without words. In C. Blakemore and S. Greenfield (Eds), *Mindwaves* (pp. 123–137). Oxford: Blackwell.
- Tellegen, A., and Atkinson, G.** (1974). Openness to absorbing and self-altering experiences ('absorption'), a trait related to hypnotic susceptibility. *Journal of Abnormal Psychology*, 83(3), 268–277.
- Thiele, A., and Stoner, G.** (2003). Neuronal synchrony does not correlate with motion coherence in cortical area MT. *Nature*, 421(6921), 366–370.
- Thomas, J. W., and Cohen, M.** (2014). A methodological review of meditation research. *Frontiers in Psychiatry*, 5, 74.
- Thomas, N.J.T.** (2016). Mental imagery. In E. N. Zalta (Ed.), *The Stanford encyclopedia of philosophy*, Summer 2016 Edition. <http://plato.stanford.edu/archives/sum2016/entries/mental-imagery/>
- Thompson, E.** (Ed.) (2001). Between ourselves: Second-person issues in the study of consciousness. Special issue, *Journal of Consciousness Studies*, 8(5–7). Also (1999) Thorverton, Devon: Imprint Academic.
- Thompson, E.** (2014). *Waking, dreaming, being: Self and consciousness in neuroscience, meditation, and philosophy*. New York: Columbia University Press.
- Thompson, E., and Zahavi, D.** (2007). Philosophical issues: Phenomenology. In P. D. Zelazo, M. Moskovitch, and E. Thompson (Eds), *The Cambridge handbook of consciousness* (pp. 67–87). Cambridge: Cambridge University Press.
- Thompson, K. G., Biscoe, K. L., and Sato, T. R.** (2005). Neuronal basis of covert spatial attention in the frontal eye field. *The Journal of Neuroscience*, 25(41), 9479–9487.
- Thomson, D. R., Besner, D., and Smilek, D.** (2013). In pursuit of off-task thought: Mind wandering-performance trade-offs while reading aloud and color naming. *Frontiers in Psychology*, 4, article 360.
- Tolstoy, L.** (1869). *War and peace* [Война́ и мир]. Full text available at <https://www.gutenberg.org/files/2600/2600-h/2600-h.htm> (trans. L. and A. Maude); also <http://ilibrary.ru/text/11/p.1/index.html> (original Russian)
- Tomasello, M.** (1999). *The cultural origins of human cognition*. Cambridge, MA: Harvard University Press.

- Tononi, G.** (2004). An information integration theory of consciousness. *BMC Neuroscience*, 5, 42.
- Tononi, G.** (2007). The information integration theory of consciousness. In M. Velmans and S. Schneider (Eds), *The Blackwell companion to consciousness* (pp. 287–299). Oxford: Blackwell.
- Tononi, G.** (2008). Consciousness as integrated information: A provisional manifesto. *The Biological Bulletin*, 215(3), 216–242.
- Tononi, G.** (2015). Integrated information theory. *Scholarpedia*, 10(1), 464. www.scholarpedia.org/article/Integrated_information_theory
- Tononi, G., and Cirelli, C.** (2003). Sleep and synaptic homeostasis: A hypothesis. *Brain Research Bulletin*, 62(2), 143–150.
- Tononi, G., and Edelman, G.** (1998). Consciousness and complexity. *Science*, 282, 5395.
- Tononi, G., and Koch, C.** (2008). The neural correlates of consciousness: An update. *Annals of the New York Academy of Sciences*, 1124, 239–261.
- Tononi, G., and Koch, C.** (2015). Consciousness: Here, there and everywhere? *Philosophical Transactions of the Royal Society B*, 370(1668), 20140167.
- Tooby, J., and Cosmides, L.** (2005). Conceptual foundations of evolutionary psychology. In D. M. Buss (Ed.), *The handbook of evolutionary psychology* (pp. 5–67). Hoboken, NJ: Wiley.
- Trapp, S., Schroll, H., and Hamker, F.H.** (2012). Open and closed loops: A computational approach to attention and consciousness. *Advances in Cognitive Psychology*, 8(1), 1–8.
- Treisman, A.** (2003). Consciousness and perceptual binding. In A. Cleeremans (Ed.), *The unity of consciousness: Binding, integration and dissociation* (pp. 95–113). New York: Oxford University Press.
- Treisman, A., and Gelade, G.** (1980). A feature integration theory of attention. *Cognitive Psychology*, 12, 97–136.
- Trent-von Haesler, N., and Beauregard, M.** (2013). Near-death experiences in cardiac arrest: Implications for the concept of non-local mind. *Archives of Clinical Psychiatry (São Paulo)*, 40(5), 197–202.
- Triplett, N.** (1900). The psychology of conjuring deceptions. *American Journal of Psychology*, 11, 439–510.
- Troscianko, E.T.** (2012). Dying by inches. In C.W. LeCroy and J. Holschuh (Eds), *First-person accounts of mental illness and recovery* (pp. 239–262). Hoboken, NJ: Wiley.
- Troscianko, E.T.** (2014). *Kafka's cognitive realism*. New York: Routledge.

Troscianko, J., and Rutz, C. (2015). Activity profiles and hook-tool use of New Caledonian crows recorded by bird-borne video cameras. *Biology Letters*, 11(12), 20150777.

Troscianko, J., Von Bayern, A. M., Chappell, J., Rutz, C., and Martin, G. R. (2012). Extreme binocular vision and a straight bill facilitate tool use in New Caledonian crows. *Nature Communications*, 3, 1110.

Tsakiris, M., and Haggard, P. (2005). The rubber hand illusion revisited: Visuotactile integration and self-attribution. *Journal of Experimental Psychology: Human Perception and Performance*, 31, 80–91.

Turing, A. (1950). Computing machinery and intelligence. *Mind*, 59, 433–460. Reprinted in J. Haugeland, (Ed.) (1997), *Mind design II: Philosophy, psychology, artificial intelligence* (pp. 29–56). Also in D. R. Hofstadter and D. C. Dennett (Eds) (1981), *The mind's I: Fantasies and reflections on self and soul* (pp. 53–67). London: Penguin.

Tuszynski, J. A. (Ed.) (2006). *The emerging physics of consciousness*. Berlin: Springer.

Tye, M. (2003). *Consciousness and persons: Unity and identity*. Cambridge, MA: MIT Press.

Uleman, J. S., Blader, S. L., and Todorov, A. (2005). Implicit impressions. In R. R. Hassin, J. S. Uleman and J. A. Bargh (Eds), *The new unconscious* (pp. 362–392). Oxford: Oxford University Press.

Ungerleider, L. G., and Mishkin, M. (1982). Two cortical visual systems. In D. J. Ingle, M. A. Goodale, and R. J. W. Mansfield (Eds), *Analysis of visual behaviour* (pp. 549–586). Cambridge, MA: MIT Press.

Utts, J. (1995). An assessment of the evidence for psychic functioning. *Journal of Parapsychology*, 59, 289–320. Reprinted in K. R. Rao (Ed.), *Basic research in parapsychology* (pp. 110–147). 2nd ed. Jefferson, NC: McFarland.

Vaitl, D., Birbaumer, N., Gruzelier, J., Jamieson, G. A., Kotchoubey, B., Kübler, A., . . . and Weiss, T. (2005). Psychobiology of altered states of consciousness. *Psychological Bulletin*, 131(1), 98–127.

van de Laar, T. (2008). Mind the methodology: Comparing heterophenomenology and neurophenomenology as methodologies for the scientific study of consciousness. *Theory & Psychology*, 18(3), 365–379.

van Eeden, F. (1913). A study of dreams. *Proceedings of the Society for Psychical Research*, 26, 431–461.

van Gulick, R. (2007). Functionalism and qualia. In M. Velmans and S. Schneider (Eds), *The Blackwell companion to consciousness* (pp. 381–395). Oxford: Blackwell.

van Laer, T., de Ruyter, K., Visconti, L. M., and Wetzels, M. (2014). The extended transportation-imagery model: A meta-analysis of the antecedents and consequences of consumers' narrative transportation. *Journal of Consumer Research*, 40(5), 797–817.

van Lommel, P. (2006). Near-death experience, consciousness, and the brain: A new concept about the continuity of our consciousness based on recent scientific research on near-death experience in survivors of cardiac arrest. *World Futures*, 62(1–2), 134–151.

van Lommel, P. (2009). Endless consciousness: A concept based on scientific studies on near-death experience. In C. D. Murray (Ed.), *Psychological scientific perspectives on out-of-body and near-death experiences* (pp. 171–186). New York: Buffalo.

van Lommel, P. (2013). Non-local consciousness: A concept based on scientific research on near-death experiences during cardiac arrest. *Journal of Consciousness Studies*, 20(1–2), 7–48.

van Lommel, P., van Wees, R., Meyers, V., and Elfferich, I. (2001). Near-death experience in survivors of cardiac arrest: A prospective study in the Netherlands. *The Lancet*, 358, 2039–2045.

VanRullen, R., Carlson, T., and Cavanagh, P. (2007). The blinking spotlight of attention. *Proceedings of the National Academy of Sciences of the United States of America*, 104(49), 19204–19209.

Varela, F. J. (1996). Neurophenomenology: A methodological remedy for the hard problem. *Journal of Consciousness Studies*, 3(4), 330–349. Also in J. Shear (Ed.) (1997). *Explaining consciousness: The 'hard problem'* (pp. 337–357). Cambridge, MA: MIT Press.

Varela, F. J. (1999). Steps to a science of inter-being: Unfolding the dharma implicit in modern cognitive science. In G. Watson, S. Batchelor, and G. Claxton (Eds), *The psychology of awakening* (pp. 71–89). London: Random House.

Varela, F. J., and Shear, J. (Eds) (1999). The view from within: First-person approaches to the study of consciousness. Special issue, *Journal of Consciousness Studies*, 6(2–3). Also (1999) Thorverton: Imprint Academic.

Varela, F. J., Thompson, E. and Rosch, E. (1991). *The embodied mind*. London: MIT Press.

Varela, F. J. (2001). Intimate distances: Fragments for a phenomenology of organ transplantation. *Journal of Consciousness Studies*, 8(5–7), 259–271. Also at www.oikos.org/varelafragments.htm

Veenendaal, M. V., Painter, R. C., Rooij, S. R., Bossuyt, P. M., Post, J. A. M., Gluckman, P. D., . . . and Roseboom, T. J. (2013). Transgenerational effects of prenatal exposure to the 1944–1945 Dutch famine. *BJOG: An International Journal of Obstetrics & Gynaecology*, 120(5), 548–554.

Velmans, M. (1999). Intersubjective science. *Journal of Consciousness Studies*, 6(2–3), 29–306. Also in F. J. Varela and J. Shear (Eds) (1999). *The view from within: First-person approaches to the study of consciousness* (pp. 299–306). Thorverton, Devon: Imprint Academic.

Velmans, M. (2000). *Understanding consciousness*. London: Routledge.

• R E F E R E N C E S

Velmans, M. (2009). *Understanding consciousness*. 2nd ed. London: Routledge.

Vimal, R.L.P. (2009). Meanings attributed to the term 'consciousness': An overview. *Journal of Consciousness Studies*, 16(5), 9–27.

Vohs, K.D., and Schooler, J.W. (2008). The value of believing in free will: Encouraging a belief in determinism increases cheating. *Psychological Science*, 19(1), 49–54.

Volkow, N.D., Baler, R.D., Compton, W.M., and Weiss, S.R.B. (2014). Adverse health effects of marijuana use. *New England Journal of Medicine*, 370(23), 2219–2227.

Vollenweider, F.X., and Kometer, M. (2010). The neurobiology of psychedelic drugs: Implications for the treatment of mood disorders. *Nature Reviews Neuroscience*, 11, 642–651.

von Hoffmannsthal, H. (1901). Letters of the returning one [Die Briefe des Zurückgekehrten], IV. First book edition (1907): *Die prosaischen Schriften*, vol. 3. Berlin: S. Fischer. Full text available at <http://gutenberg.spiegel.de/buch/die-briefe-des-zuruckgekehrten-987/4> (original German)

Vorberg, D., Mattler, U., Heinecke, A., Schmidt, T., Schwarzbach, J. (2003). Different time courses for visual perception and action priming. *Proceedings of the National Academy of Sciences, USA*, 100(1), 6275–6280.

Vorhees, B. (2000). Dennett and the deep blue sea. *Journal of Consciousness Studies*, 7(3), 53–69.

Voss, U., Schermelleh-Engel, K., Windt, J., Frenzel, C., Hobson, A. (2013). Measuring consciousness in dreams: The lucidity and consciousness in dreams scale. *Consciousness and Cognition*, 22, 8–21.

Voss, U., Holzmann, R., Tuin, I., and Hobson, J.A. (2009). Lucid dreaming: A state of consciousness with features of both waking and non-lucid dreaming. *Sleep*, 32, 1191–1200.

Voss, U., Holzmann, R., Hobson, A., Paulus, W., Koppehele-Gossel, J., and Klimke, A. (2014). Induction of self-awareness in dreams through frontal low current stimulation of gamma activity. *Nature Neuroscience*, 17(6), 810–812.

Vossel, S., Geng, J.J., and Fink, G.R. (2014). Dorsal and ventral attention systems: Distinct neural circuits but collaborative roles. *The Neuroscientist*, 20(2), 150–159.

Wade, K.A., Garry, M., Read, J.D., and Lindsay, S. (2002). A picture is worth a thousand lies: Using false photographs to create false childhood memories. *Psychonomic Bulletin & Review*, 9(3), 597–603.

Wagstaff, G. (1994). Hypnosis. In A.M. Colman (Ed.), *Companion encyclopedia of psychology*, vol. 2 (pp. 991–1006). London: Routledge.

- Wallace, B., and Fisher, L. E.** (1991). *Consciousness and behavior*. 3rd ed. Boston, MA: Allyn and Bacon.
- Wamsley, E. J., Peery, K., Djonlagic, I., Reaven, L. B., and Stickgold, R.** (2010). Cognitive replay of visuomotor learning at sleep onset: Temporal dynamics and relationship to task performance. *SLEEP*, 33(1), 59–68.
- Wamsley, E. J.** (2014). Dreaming and offline memory consolidation. *Current Neurology and Neuroscience Reports*, 14(3), 433.
- Wamsley, E. J., and Stickgold, R.** (2011). Memory, sleep and dreaming: Experiencing consolidation. *Sleep Medicine Clinics*, 6(1), 97–108.
- Ward, J.** (2013). Synesthesia. *Annual Review of Psychology*, 64, 49–75.
- Ward, J., Huckstep, B., and Tsakanikos, E.** (2006). Sound-colour synaesthesia: To what extent does it use cross-modal mechanisms common to us all? *Cortex*, 42(2), 264–280.
- Ward, J., and Wright, T.** (2014). Sensory substitution as an artificially acquired synaesthesia. *Neuroscience and Biobehavioral Reviews*, 41, 26–35.
- Watanabe, M., Cheng, K., Murayama, Y., Ueno, K., Asamizuya, T., Tanaka, K., and Logothetis, N.** (2011). Attention but not awareness modulates the BOLD signal in the human V1 during binocular suppression. *Science*, 334, 829–831.
- Waters, F.A.V., Badcock, J. C., and Maybery, M. T.** (2003). Revision of the factor structure of the Launay-Slade Hallucination Scale (LSHS-R). *Personality and Individual Differences*, 35(6), 1351–1357.
- Waters, F., Collerton, D., ffychte, D.H., Jardri, R., Pins, D., Dudley, R., Blom, J.D., Mosimann, U.P., Eperjesi, F., Ford, S., and Larøi, F.** (2014). Visual hallucinations in the psychosis spectrum and comparative information from neurodegenerative disorders and eye disease. *Schizophrenia Bulletin*, 40(4), 5233–5245.
- Watson, G., Batchelor, S., and Claxton, G.** (Eds) (1999). *The psychology of awakening: Buddhism, science and our day-to-day lives*. London: Rider.
- Watson, J. B.** (1913). Psychology as the behaviorist views it. *Psychological Review*, 20(2), 158–177. Reprinted (1994) in *Psychological Review*, 101(2), 248–253.
- Watts, A. W.** (1957). *The way of Zen*. New York: Pantheon Books.
- Watts, A.** (1961). *Psychotherapy East and West*. London: Jonathan Cape.
- Watts, P.** (2006). *Blindsight*. New York: Tom Doherty, Tor.
- Watzl, S.** (2011). The nature of attention. *Philosophical Compass*, 6(11), 842–853.
- Watzl, S.** (2017). *Structuring mind: The nature of attention and how it shapes consciousness*. Oxford: Oxford University Press.

• R E F E R E N C E S

- Webb, T.W., and Graziano, M.S.A.** (2015). The attention schema theory: A mechanistic account of subjective awareness. *Frontiers in Psychology*, 6, article 500.
- Webster, R.** (1995). *Why Freud was wrong: Sin, science and psychoanalysis*. London: HarperCollins.
- Wegner, D.M.** (1989). *White bears and other unwanted thoughts: Suppression, obsession, and the psychology of mental control*. New York: Viking Penguin.
- Wegner, D.M.** (2002). *The illusion of conscious will*. Cambridge, MA: MIT Press.
- Wegner, D. M.** (2003). The mind's best trick: How we experience conscious will. *Trends in Cognitive Sciences*, 7(2), 65–69.
- Wegner, D.** (2005). Who is the controller of controlled processes? In R. R. Hassin, J. S. Uleman and J. A. Bargh (Eds), *The new unconscious* (pp. 19–36). Oxford: Oxford University Press.
- Wegner, D.M., and Wheatley, T.** (1999). Apparent mental causation: Sources of the experience of will. *American Psychologist*, 54, 480–492.
- Wehrle, R., Kaufmann, C., Wetter, T.C., Holsboer, F., Auer, D.P., Polimächer, T., and Czisch, M.** (2007). Functional microstates within human REM sleep: First evidence from fMRI of a thalamocortical network specific for phasic REM periods. *European Journal of Neuroscience*, 25, 863–871.
- Weil, A.** (1998). *The natural mind: A new way of looking at drugs and the higher consciousness*. Rev. ed. New York: Houghton Mifflin.
- Weir, A.A.S., Chappell, J., and Kacelnik, A.** (2002). Shaping of hooks in New Caledonian crows. *Science*, 297(5583), 981.
- Wei Wu Wei.** (2004). *Open secret*. Hong Kong: Hong Kong University Press.
- Weiskrantz, L.** (1986). *Blindsight: A case study and implications*. Oxford: Oxford University Press.
- Weiskrantz, L.** (1997). *Consciousness lost and found*. Oxford: Oxford University Press.
- Weiskrantz, L.** (2007). The case of blindsight. In M. Velmans and S. Schneider (Eds), *The Blackwell companion to consciousness* (pp. 175–180). Oxford: Blackwell.
- Weizenbaum, J.** (1966). ELIZA – a computer program for the study of natural language communication between man and machine. *Communications of the ACM*, 9(1), 36–45.
- Welsh, T.** (2015). It feels instantaneous, but how long does it really take to think a thought? *The Conversation*, 26 June 2015. <http://theconversation.com/it-feels-instantaneous-but-how-long-does-it-really-take-to-think-a-thought-42392>

- West, D. J.** (1948). A mass observation questionnaire on hallucinations. *Journal of the Society for Psychical Research*, *34*, 187–196.
- West, L. J.** (1962). *Hallucinations*. New York: Grune & Stratton.
- West, M. A.** (Ed.) (1987). *The psychology of meditation*. Oxford: Clarendon Press.
- Wheeler, M., and Clark, A.** (2008). Culture, embodiment and genes: Unravelling the triple helix. *Philosophical Transactions of the Royal Society B*, *363*, 3563–3575.
- Whiten, A., and Byrne, R. W.** (1997). *Machiavellian intelligence II: Extensions and evaluations*. Cambridge: Cambridge University Press.
- Whitwell, R. L., Milner, A. D., and Goodale, M. A.** (2014). The two visual systems hypothesis: New challenges and insights from visual form agnostic patient DF. *Frontiers in Neurology*, *5*, article 255.
- Wilber, K.** (1997). An integral theory of consciousness. *Journal of Consciousness Studies*, *4*(1), 71–92.
- Wilber, K.** (2001). *A brief history of everything*. Boston, MA: Shambhala.
- Wilber, K., Engler, J., and Brown, D.** (Eds) (1986). *Transformations of consciousness: Conventional and contemplative perspectives on development*. Boston, MA: Shambhala.
- Wilkins, L. K., Girard, T. A., and Cheyne, J. A.** (2011). Ketamine as a primary predictor of out-of-body experiences associated with multiple substance use. *Consciousness and Cognition*, *20*(3), 943–950.
- Wilkins, L. K., Girard, T. A., and Cheyne, J. A.** (2012). Anomalous bodily-self experiences among recreational ketamine users. *Cognitive Neuropsychiatry*, *17*(5), 415–430.
- Wilkinson, S.** (2014). Accounting for the phenomenology and varieties of auditory verbal hallucination within a predictive processing framework. *Consciousness and Cognition*, *30*, 142–155.
- Williams, B. J.** (2011). Revisiting the ganzfeld ESP debate: a basic review and assessment. *Journal of Scientific Exploration*, *25*(4), 639–661.
- Williams, G. C.** (1966). *Adaptation and natural selection*. Princeton, NJ: Princeton University Press.
- Williams, M. A., Morris, A. P., McGlone, F., Abbott, D. F., and Mattingley, J. B.** (2004). Amygdala responses to fearful and happy facial expressions under conditions of binocular suppression. *Journal of Neuroscience*, *24*(12), 2898–2904.
- Wilson, A.** (2012). Patient DF uses haptics, not intact visual perception-for-action to reach for objects. *Notes from Two Scientific Psychologists*, 13

• R E F E R E N C E S

April 2012. <http://psychsciencenotes.blogspot.co.uk/2012/04/patient-df-uses-haptics-not-intact.html>

Wilson, B.A., and Wearing, D. (1995). Prisoner of consciousness: A state of just awakening following herpes simplex encephalitis. In R. Campbell and M. Conway (Eds), *Broken memories: Case studies in memory impairment* (pp. 14–30). Oxford: Blackwell.

Wilson, D.S., and Sober, E. (1994). Reintroducing group selection to the human behavioral sciences. *Behavioral and Brain Sciences*, 17(4), 585–654 (incl. commentaries and authors' response).

Wilson, D. S., and Wilson, E. O. (2008). Evolution 'for the good of the group'. *American Scientist*, 96(5), 380–389.

Wilson, E. O. (1975). *Sociobiology: The new synthesis*. Cambridge, MA: Harvard University Press.

Wimsatt, W. (2010). Memetics does not provide a useful way of understanding cultural evolution. In F. Ayala and R. Arp (Eds), *Contemporary debates in philosophy of biology* (pp. 273–291). Chichester: Wiley-Blackwell.

Windey, B., Gevers, W., and Cleeremans, W. (2013). Subjective visibility depends on level of processing. *Cognition*, 129(2), 404–409.

Windt, J.M., and Noreika, V. (2011). How to integrate dreaming into a general theory of consciousness – a critical review of existing positions and suggestions for future research. *Consciousness and Cognition*, 20(4), 1091–1107.

Winfield, A.F., and Griffiths, F. (2010). Towards the emergence of artificial culture in collective robot systems. In P. Levi and S. Kernbach (Eds), *Symbiotic multi-robot organisms: Reliability, adaptability, evolution* (pp. 425–433). Berlin: Springer.

Winfield, A.F. (2017). When robots tell each other stories: The emergence of artificial fiction. In R. Walsh and S. Stepney (Eds), *Narrating complexity* (in press). London: Springer.

Winkelman, M. (2014). Psychedelics as medicines for substance abuse rehabilitation: Evaluating treatments with LSD, peyote, ibogaine and ayahuasca. *Current Drug Abuse Reviews*, 7(2), 101–116.

Wiseman, R., Greening, E., and Smith, M. (2003). Belief in the paranormal and suggestion in the seance room. *British Journal of Psychology*, 94, 285–297.


Wiseman, R., and Milton, J. (1998). Experiment One of the SAIC remote viewing program: A critical re-evaluation. *Journal of Parapsychology*, 62, 297–308.

Wiseman, R., Smith, M., and Kornbrot, D. (1996). Exploring possible sender-to-experimenter acoustic leakage in the PRL autoganzfeld experiments. *Journal of Parapsychology*, 60, 97–128.

- Wong, J. I., and Sonnad, N.** (2016). Google won the game Go by defying millennia of basic human instinct. *Quartz*, 25 March. <https://qz.com/639952/googles-ai-won-the-game-go-by-defying-millennia-of-basic-human-instinct/>
- Woolf, V.** (1915). Woolf, V. (1915). *The voyage out*. London: Duckworth. Full text available at <http://www.gutenberg.org/ebooks/144>
- Wren-Lewis, J.** (1988). The darkness of God: A personal report on consciousness transformation through an encounter with death. *Journal of Humanistic Psychology*, 28, 105–122.
- Wren-Lewis, J.** (2004). Personal communication. In conversation with Sue, 19 August.
- Wright, E. L.** (Ed.) (2008). *The case for qualia*. Cambridge, MA: MIT Press.
- Wu, W.** (2011). Attention as selection for action. In C. Mole, D. Smithies, and W. Wu (Eds), *Attention: Philosophical and psychological essays* (pp. 97–116). New York: Oxford University Press.
- Wundt, W. M.** (1897). *Outlines of psychology*. Trans. C. H. Judd. Leipzig: Engelmann.
- Xu, J., Vik, A., Groote, I. R., Lagopoulos, J., Holen, A., Ellingen, Ø., Håberg, A. K., and Davanger, S.** (2014). Nondirective meditation activates default mode network and areas associated with memory retrieval and emotional processing. *Frontiers in Human Neuroscience*, 8, article 86.
- Yeshurun, Y., and Carrasco, M.** (1998). Attention improves or impairs visual performance by enhancing spatial resolution. *Nature*, 396, 72–75.
- Young, A. W.** (1996). Dissociable aspects of consciousness. In M. Velmans (Ed.), *The science of consciousness* (pp. 118–139). London: Routledge.
- Zahavi, D.** (2011). Unity of consciousness and the problem of self. In S. Gallagher (Ed.), *The Oxford handbook of the self* (pp. 316–335). New York: Oxford University Press.
- Zeki, S.** (2001). Localization and globalization in conscious vision. *Annual Review of Neuroscience*, 24, 57–86.
- Zeki, S.** (2003). The disunity of consciousness. *Trends in Cognitive Sciences*, 7(5), 214–218.
- Zeki, S.** (2007). A theory of micro-consciousness. In M. Velmans and S. Schneider (Eds), *The Blackwell companion to consciousness* (pp. 580–588). Oxford: Blackwell.
- Zeki, S.** (2008). The disunity of consciousness. *Trends in Cognitive Sciences*, 7(5), 214–218.
- Zeki, S.** (2015). Area V5 – a microcosm of the visual brain. *Frontiers in Integrative Neuroscience*, 9, article 21.

• R E F E R E N C E S

- Zeki, S., and Bartels, A.** (1998). The asynchrony of consciousness. *Proceedings of the Royal Society B*, 265, 1583–1585.
- Zeki, S., and Bartels, A.** (1999). Toward a theory of visual consciousness. *Consciousness and Cognition*, 8, 225–259.
- Zeki, S., and Marini, L.** (1998). Three cortical stages of colour processing in the human brain. *Brain*, 121, 1669–1685.
- Zeman, A.** (2001). Consciousness. *Brain*, 124, 1263–1289.
- Zentall, T.R.** (2006). Imitation: Definitions, evidence, and mechanisms. *Animal Cognition*, 9(4), 335–353.
- Zhuangzi.** (3rd century BC/2010/2016). *Zhuangzi: The inner chapters*. Trans. R. Eno. <http://www.indiana.edu/~p374/Zhuangzi.pdf>
- Ziat, M., Smith, E., Brown, C., DeWolfe, C., and Hayward, V.** (2014, February). Ebbinghaus illusion in the tactile modality. *2014 IEEE Haptics Symposium (HAPTICS)*, 581–585.
- Zohar, D., and Marshall, I.** (2002). *The quantum soul*. London: HarperCollins.



Index

- 4E cognition (embodied, embedded, extended, enactive) 26, 311; *see also* embodied, cognition; enactive cognition
- Aaronson, Scott 48, 140–141
- abdominal breathing 179
- Abhidharma* 494
- 'aboutness' 23
- absorption 212, 349, 355, 392, 422
- AC *see* artificial consciousness
- ACC *see* anterior cingulate cortex
- access consciousness 34, 43, 88, 136, 168, 172–173, 477; *see also* phenomenal consciousness
- acetylcholine 348
- acid trip *see* LSD
- A-consciousness *see* access consciousness
- action: conscious vs. unconscious 194–197
- adaptation 196, 254–255, 258, 287, 296, 343
- Adaptation and Natural Selection* 255
- addiction 184, 354, 361
- Adolphs, Ralph 87, 100
- Advaita 493
- afterimages 62, 393
- afterlife hypothesis 428, 437
- agency 218–245; attribution of 238, 379; brain location of 429; independent 387; invisible 121; v reception 294; and self 453, 457; sense of 29, 184, 189, 220, 237, 242, 420
- aggression 242, 277, 279, 400
- Aglioti, Salvatore 204
- agnosia 169, 203, 208
- AI *see* artificial intelligence
- AIM model 348, 349, 402, 405
- Albahari, Miri 455
- alcohol 21, 150, 344–355, 357, 393
- Aleksander, Igor 308, 329, 330, 331, 503
- algorithm 46, 121, 213, 222, 296; computational 306–307, 324, 331, 384–385; evolutionary 213, 251, 298, 310, 324; Google search 222
- alien 237, 274, 379, 413, 450; abduction 382, 396, 414–415; hand 222
- Alkire, Michael T. 93
- Alpert, Richard 361
- alphaGo 315–316
- alpha rhythm 79, 182
- altered states of consciousness (ASCs) 343–371; definition of 344–345, 366–368; drugs and altered states 25, 179, 351–362, 368, 423; hypnosis as 365–366, 369–370, 390; mapping 347–350, 402–403; meditation and 362–364, 368; mental illness and 367–368; self in 263; sleep, and dreams and 351, 370, 397, 398, 408; TMS induced 80
- altruism 277, 279
- Alzheimer's disease 150, 338
- American Sign Language (ASL) 264, 272
- amnesia 129, 143, 149, 150, 191, 449; anterograde 149, 150; dense 150; retrograde 149; source 94
- amodal perception 58
- amoeba 141, 257, 293
- amphetamines 347, 351–352, 383, 384
- amputee 98–99
- aMuu 338
- amygdala 80, 193, 399, 401, 406
- anaesthesia 93, 121, 191–192, 349, 353
- anaesthetic 93–94, 191–192, 214, 215, 264, 352–354, 423; isoflurane 93; ketamine 93, 353–354, 423; laughing gas 93, 352; nitrous oxide 93, 352–353; Propofol 93, 94

- Analytical Engine 305, 323
 anarchic hand 223
anatta 244, 436, 439, 505
 Andrade, Jackie 36, 277
 angels 252, 320, 430
 animals: amoeba 141, 257, 293; ape 261; bat 32–34, 44, 279, 332, 461, 477; butterfly 92, 177, 256, 267, 492; chimpanzees 146, 256, 261, 264–272, 279, 289; cockroach 20; consciousness 259, 262, 326; corvids 265; crows 261, 265, 274; dog 45, 263, 269, 338, 405, 461; evolution and animal minds 249–275; frog 258–259, 425; gorilla 66, 169, 256, 264–265, 272; horse 23, 32, 256; imitation 265, 269–272, 282; language 82, 257, 264, 272–273; learning 24, 261, 270; lobster 257, 261–262; monkey 61, 82, 88, 91, 256, 264–267; octopus 257–258, 260–261, 273–274; orangutan 256, 263–264, 272; parrot 265, 272, 484; pigeon 24, 244, 265; rabbit 153–154, 239, 284, 391; rat 24, 137, 254–256, 405, 406, 430; raven 261, 269; salmon 80; self-recognition 262–267; snakes 152, 256, 377, 385, 394; suffering 261–262, 273; termites 261, 270, 312; theory of mind 267–269, 273, 285; worms 32, 260
 animation 330
 ANNs *see* artificial neural network
 anorexia 365–366
 anoxia 430
 anterior cingulate 116, 162, 222–223, 232, 366, 399, 406
 anterior cingulate cortex (ACC) 86, 95–97, 181, 480
 anterior commissure 81, 89
 anterograde amnesia 149–150
 anxiety 146, 176, 211–212, 352, 358, 363–365
 apparitions 413
 archetype 382, 401
 argument from design 249–250
 Aristotle 303, 447
 arousal 79–80, 134, 346–350, 380, 416, 418–419
 artificial consciousness (AC) 308, 318, 320–322, 325, 330; building a conscious machine 328–331; strong AC 308; weak AC 308
 artificial intelligence (AI) 303–308; good old-fashioned (GOFAI) 307, 309, 311, 313, 326, 335; moderate AI 327; rule-and-symbol 307, 310; strong AI 310, 318, 326; weak AI 308
 artificial neural network (ANN) 26, 309, 310, 384
 artificial scotoma 61
 Aru, Jaarn 87
 ASCs *see* altered states of consciousness
 as-if: agency 332, 458, 507; consciousness 308, 319, 329, 339, 485, 507; creativity 324; free will 242–5, 507; intentionality 267, 304–305, 326–329, 485, 507; self 458, 507; thinking 318
 Asimo (Honda’s humanoid robot) 323
 ASL *see* American Sign Language
 assisted suicide 279
 ‘Astonishing’ Hypothesis 118, 160, 241
 astral body 423, 424, 425, 461
 astral projection 417, 422, 423
 atheist 493
atman 493, 505
Atropa belladonna 394
 atropine 293
 attention 159–184; and binding 132, 166, 171; and consciousness 159–163, 168–174; covert 164, 167; as easy problem 46; focused 69, 132, 160, 177–178, 182, 347; interoceptive 183; involuntary 161, 162; and meditation 175–184, 362–363; and memory 166; schema theory 167, 171, 290, 295, 333, 453; selective 134, 170, 181, 353; spatial 166, 167; spotlight of 113–114, 159–163, 165, 167, 172; training of 466
 attentional blink 90, 116, 163–164
 auditory cortex 78, 86, 210, 399, 430
 aura 379, 395
 Austin, James 178, 181, 364, 495, 502
 autism 144, 162, 211
 autobiographical: memory 114, 265, 375, 450; self 265, 449, 450
 automata 19, 305; Cartesian 261; Conscious 19, 200, 257, 305
 automatisaion 95, 179, 194
 automaton theory 28, 200
 autopoiesis 473
 autoscopy 421, 422, 425
 awareness: self 119, 181, 257, 263, 265, 350; visual 84, 90, 118, 132, 153, 205
 awareness neurons 57, 118, 154
 AWARE project 429
 ayahuasca 355, 359, 361, 394, 395
 Baars, Bernard 95, 113–116, 119, 196, 257, 258, 285, 331, 450, 451, 477, 478, 503; Profile 113; *see also* Global Workspace Theory
 Baba Ram Dass 361

- Babbage, Charles 305, 323
 Bachmann, Talis 87
 Bach-y-Rita, Paul 208
 back-propagation algorithm 310
 backwards: masking 155; referral 226–227, 232
 Balzac, Honoré 355
 barbiturates 383
 Barlow, Horace 292, 293
 bat 32–34, 44, 279, 332, 461, 477
 Batchelor, Stephen 175, 493, 497
 battle of As and Bs 467–471
 Bauby, Jean-Dominique 92
 Baudelaire, Charles 355
 Baxter, Stephen 274
 Bayne, Tim 148, 232
 Beaton, Michael 39, 40
 Beauchamp, Christine 440–442
 behaviourism 17, 20, 24–25, 28, 108, 198, 289, 387, 487
Being No One 350
 belladonna 394
 benign user illusion 300, 334
 Bennett, Maxwell R. 147, 148
 Bentall, Richard 377, 379
 Berger, Hans 79
 Berkeley, George 18
 β -carbolines 359
 bicameral mind 285–286
 binding 119, 120, 136, 407; intentional 233; problem 120, 129, 130–132, 166; by synchrony 132–135, 171; temporal 132; visual 46, 131, 133
 binocular: rivalry 62, 82, 88, 90–91, 386; suppression 170
 biological naturalism 322
 bipolar 366
 Bisiach, Eduardo 151
 Blackmore, Susan 3, 63, 71, 72, 84, 336, 382, 391, 410, 411, 414; profile 2
 Blakemore, Sarah 232
 Blakeslee, Sandra 58, 59, 61, 63, 205
 Blanke, Olaf 426, 428
 blind: spots 58–62, 63, 373, 479; variation 251, 299
 blindness 86, 381; change 54, 63–65, 66, 67, 68–69, 71, 259, 313; inattentional 54, 65–68; partial 381
 blindsight 205, 206–214, 449; agnosia and 188; and brain damage 86, 136, 151
 blindsmell 206
 Block, Ned 34, 35, 88, 136, 160, 172, 173, 198, 209
 blood oxygen level dependent (BOLD) 80, 428
 Bloom, Paul 437
 body image 93, 346, 424, 425, 428, 450
 body language 67, 269, 337
 body-swap illusion 263, 481
 Boole, George 306
 Bourne, Rev. Ansel 439–440
Brahman 493
 brain 26, 77–81; damage 61, 136, 203–204, 210, 222, 381, 422, 449; -imaging 79–80, 224, 361; metabolism 49, 79, 80, 93; scanning 79, 80, 85, 115, 425, 459; size 259, 265, 267, 284; stem 92, 93, 96, 99, 179, 399, 450
 brain lobes: frontal 81, 135, 138, 212, 404; occipital 78, 81, 86, 162; parietal 78, 81, 86, 162, 416, 420; temporal 61, 81, 413, 424, 430
 brain–mind space 348, 349, 402
 brains and computers 308–310
 brainstem 92, 93, 96, 99, 179, 399, 450
 brain-teasers 212
 Braithwaite, Jason 425, 430
 Brasington, Leigh 363
 Breazeal, Cynthia 337–338
 Breitmeyer, Bruno 90
 Brentano, Franz 23
 bridge locus 85
 Broca's area 78, 222
 Broks, Paul 445
 Brontë, Emily 399
 Brooks, Rodney 312, 313, 323, 337, 338, 459, 460
 B-Space 350
 Buddha 436–438, 490–495, 502, 505, 506
 Buddhism 436–439, 490–498, 501–506; meditation and 175, 178–180, 473; in science 492–496; Tibetan 427; Zen 180, 363, 493, 495–496; *see also anatta; jhanas; kensho*
 bundle theory 436–439, 442–448, 451–454, 457–459, 505
 Café Wall illusion 54, 55, 376
 Calvin, William 297, 478
 camouflage 267
 Campbell, Donald 251
 cannabis 344, 345, 347, 355–357, 363, 369, 383
 Carpenter, William Benjamin 195
 Carrasco, Marisa 87, 174
 Carrington, Whateley 423
 Carruthers, John 261
 Cartesian: daydreams 115; dogma of ghost in machine 16, 18; dualism 104, 122,

- 439; materialism (CM) 18, 58, 104–105, 134, 166, 192, 207, 214, 233
- Cartesian Theatre 18, 104–107, 111–113, 135, 487, 499; audience in the 122, 292, 333, 386, 407, 409, 447, 458; Kismet has no 339; pernicious 115, 407; there is no 409, 457
- cassette theory of dreams 409
- Castaneda, Carlos 394–395
- Castiello, Umberto 203
- caudate nucleus 86
- causal efficacy of consciousness 199–200, 288
- causal paradox 200
- cause and effect 16, 506
- Census of Hallucinations 377–378
- Center of Narrative Gravity 336, 456–459
- central nervous system 79, 228
- cerebellum 78–79, 150, 222, 232
- C-fibres 96–97
- Chalmers, David 3, 12, 18, 20, 28–29, 38–39, 41–50, 88, 140–141, 210, 321, 327, 330; hard problem 3, 28–29, 44–50; profile 41
- change: blindness 54, 63–65, 259, 313; detection 65, 68
- chaos theory 251, 310
- Charles Bonnet syndrome 381, 384, 386
- chasm 27, 29, 85, 183, 362, 449
- Cheesman, James 190
- Cheney, Dorothy 267
- chess 305, 306, 315–316
- Chiang, Ted 243
- chimpanzees 146, 221, 256, 261, 264–272, 279, 289, 290
- Chinese Room 308, 325, 327–330, 335
- Chopra, Deepak 120, 253
- Chrisley, Ron 258, 327, 330, 334
- Christianity 175, 180, 219, 244, 299, 437, 493, 503
- chronic fatigue syndrome 367
- Churchland, Patricia Smith 12, 39, 43, 45, 49, 117, 122, 183, 293, 468, 476, 478; profile 49
- Churchland, Paul 39, 49, 117
- Church–Turing thesis 310
- cingulate gyrus 61, 78, 80, 162
- circadian rhythms 260
- clairvoyance 388, 422–423, 430
- Clark, Andy 65, 78, 91, 233, 311, 313, 456, 468; Profile 214
- Clarke, Arthur. C. 392
- Clarke, Chris 46
- classical conditioning 24, 149
- Claxton, Guy 213, 240, 244, 295, 497–498
- Clowes, Robert 303
- CM see Cartesian, materialism
- CMF see Libet's Theories, conscious mental field
- CNS see central nervous system
- CNS depressants 79, 228
- cocaine 500
- co-consciousness 441
- CODAM model 331
- CogAff architecture 334
- cognitive: maps 80; science 26, 28, 473–474, 476, 495; unconscious 31
- Cog robot 322–324
- Cohen, David 25, 101, 144, 169, 364
- coherence 7, 46, 69, 120–122, 135, 333, 461
- Coleridge, Samuel Taylor 186, 213
- collective unconscious 382
- colour: figment 60; perception 38, 135; phi 153; qualia 38, 282, 288; scientist 38–40, 282, 328, 470; stimulus 373; vision 38–40
- coma 6, 86, 93–94, 203, 243, 349, 405
- computationalism 330
- computational theory of mind 308
- computer: analogue 308–309; digital 26, 46, 308, 309–310; digital vs. analogue 308; parallel vs. serial 309; science 25
- Conan Doyle, Arthur 164
- confabulation 123, 144, 146, 409
- connectionism 26, 309–310, 322, 327, 335
- 'conscie' 282, 284
- conscious: decision making 100, 172, 194, 214, 220; experience 20–23, 41–42, 86–88, 106, 111, 114; inessentialism 42, 209, 281, 284, 319, 470; mental field (CMF) 45, 129, 196
- consciousness: adaptive function of 279, 284, 287, 296, 404; and attention 159–163, 165, 168–174; causal efficacy of 199–200, 288; core 265, 449; in cyber space 300, 332, 460–462; defining 21, 32–33; extended 265, 449; evolution of 44, 247, 249, 280–284, 289, 298, 403; function of 44, 105, 115, 276–281, 283, 289, 470; neural basis of 81, 87, 94, 96, 260; phenomenal (P) 34–35, 50, 134–136, 148, 261, 367; in psychology 21; pure 23, 179, 472, 503–504; reflexive model of 477–480; simultaneous 441, 450; split 144; unity of 120–121, 128–155; waking consciousness 343
- contrastive analysis 95, 115, 450
- core consciousness 265, 449
- corpus callosum 78, 81, 143–145, 222

correlation and cause 81–83, 134, 241, 356
 cortex: anterior cingulate 86, 96, 162, 181, 222, 223, 401; cerebral 18, 78, 134, 297, 450; primary somatosensory 78, 86; visual (V1) 33, 38, 78, 82, 130, 361, 383
 cortical activation 237, 419
 cortico-cortical 81, 93, 119; reverberant loops 93
 cortico-thalamic loops 81, 93
 cosmic consciousness 212, 253
 Cosmides, Leda 277
 Cotterill, Rodney 141, 324
 Craik, Kenneth 307, 332
 creationism 252
 creativity 175, 187, 194, 211–214, 298, 310, 324
 CREEPI 339
 Crick, Francis 46, 57, 83–84, 85, 88, 110, 111, 118, 132, 133, 136, 154, 160, 241, 297, 404, 477, 478
 CRONOS 324, 332–333
 Crook, John 178, 288, 494, 497, 498
 cryptomnesia 94
 Csikszentmihalyi, Mihaly 212, 300, 465
 C-space 350
 CT *see* Cartesian Theatre
 CT scan *see* X-ray computed tomography
 cultural: evolution 213, 255, 270; replicator 298
 cutaneous rabbits 153
 CyberChild 324
The Cyberiad 321
 Cytowic, Richard 144

Dalai Lama 492
 Damasio, Antonio 98–99, 212, 265, 331, 449, 450, 451
 Darwin, Charles 19, 78, 250–252, 253, 254, 257, 263, 269, 297, 298, 334, 378
 Darwin, Erasmus 250, 334
Darwin's Dangerous Idea 251
 Davy, Humphrey 352
 Dawkins, Marian Stamp 261, 262
 Dawkins, Richard 33, 250, 256, 290, 296, 298, 299, 300; profile 255
 daydream 115, 182, 346, 365, 379, 392
 Deacon, Terrence 291
 deaf hearing 206
 death 175, 359, 378, 421, 423–425, 428–430, 501
 deautomization 179, 194
 deception 267, 268, 285, 483, 488, 504
 Deep Blue 315
 Deep Fritz 315
 Deep Thought 315
 default mode network (DMN) 166, 181–183, 234, 299, 353, 361, 475
 Dehaene, Stanislas 47, 87, 116, 117, 169, 173, 193, 197, 451
 Deikman, Arthur 178, 179
 Delage, Yves 413
 de la Mettrie, Julien Offray 19, 304
 delay in consciousness 224, 226
 deluded machines 333–334, 336
 delusions 367
 dementia 150, 379
 de Mille, Richard 395
 Democritus 220, 303
 Dennett, Daniel 6, 11, 18, 27, 36, 37, 39, 40, 43, 49, 50, 55, 56, 58, 59, 60, 85, 90, 104, 105, 106, 107, 108, 111, 116, 117, 122, 123, 124, 125, 126, 136, 142, 154, 155, 173, 208, 234, 235, 243, 257, 266, 282, 283, 295, 296, 297, 299, 300, 309, 313, 317, 321, 327, 386, 409, 410, 439, 449, 457, 458, 459, 466, 467, 468, 469, 470, 471, 477, 478, 483, 484, 485, 486, 487, 500; profile 104
 Depp, Johnny 459
 Depraz, Natalie 473
 depression 94, 176, 184, 344, 353, 367, 498; treating 358, 361, 363, 368
 Descartes, René 15, 17, 56, 85, 104, 105, 106, 107, 129, 195, 199, 200, 257, 290, 303, 304, 315, 316, 435, 447; profile 16
 determinism 219, 231, 241–242, 244
 dharma 490, 506
 dichotic listening 165
 DID *see* dissociative identity disorder
 Diderot 16
 digital computers 26, 308–310
 Di Mauro, Ernesto 122
 dimethyltryptamine (DMT) 359, 422
 disease 92, 255, 383, 405; Alzheimer's 150, 338; epilepsy 144; heart 176, 253, 278; iatrogenic 441; Parkinson's 150, 423
 disinhibition 383, 386, 430
 dissociation 93, 150, 201–204, 365, 422
 dissociative identity disorder (DID) 441; *see also* split personality
Diving Bell and the Butterfly 92
 DLPFC *see* dorsolateral prefrontal cortex
 DMN *see* default mode network
 DNA (deoxyribonucleic acid) 11, 32, 49, 83, 254, 255, 256, 298
 Dobzhansky, Theodosius 251
 Dogen, Eihei 363
 Donald, Merlin 291

- doorway test for auras 395
- dopamine 322, 351, 353, 359, 360, 364, 423
- Doppelgänger 421
- dorsal attention system 161–162, 166, 184
- dorsal stream 81, 131, 132, 204–205, 208, 210, 227
- dorsolateral prefrontal cortex 89, 181, 222, 223, 399, 401, 420
- double transduction 123
- DPMC 95–96
- dreams 61, 67, 114, 186, 196, 349, 398–407;
 - bizarreness 407; cassette theory of 409; day 115, 182, 346; evolution of 403–404; as experiences 398, 408; falling 412–413, 416, 425; lucid 348, 398, 405, 408, 414–422; phenomenology of 407, 409; recall 400, 404, 406; retro-selection theory of 410, 415
- dress, the 373–374
- Dr. Jekyll and Mr Hyde 440
- drugs 344–346, 350–362; amphetamines 347, 351–352, 383–384; anaesthetics 93–94, 191–192, 214–215, 352–354, 423; mind-altering 344, 346, 351; psychedelics 351, 354–369, 385, 422, 508; psychoactive 200, 350–351, 354–355, 359–360; stimulants 351–352
- dual-aspect 45–46, 479; monism 479; theory 15, 18, 45
- dualism 12, 15–18, 20, 39, 44–45, 58, 97–98, 100; Cartesian 15, 104, 120, 122, 439; Descartes 104, 106; double 232; mind-body 130; naturalistic 18; property 15, 18, 471; quantum interactive 120; substance 15–16, 129, 437, 479
- dualist theories of consciousness 15–18, 21
- dual-process theory 95, 187
- dynamic core 138, 297

- Eagleman, David. M. 153, 234
- easy problems 28–29, 44–49, 81, 100, 101
- ECCE robot 324, 332
- Eccles, John 18, 120, 129, 195, 196, 227, 231, 235, 437
- echolocation 33, 203, 273
- ecstasy 344, 351
- ectoplasm 377–378
- Edelman, Gerald M. 119, 138, 273, 296, 297, 461, 478
- EEG *see* electroencephalogram
- ego theory 439, 442–443, 456; *see also*
 - bundle theory
- Ego Tunnel, The* 350
- Ego tunnel 350, 453

- eightfold path 491
- Einstein, Albert 37, 78, 384, 429
- élan vital 11–12, 284
- electroencephalogram 79, 400
- eliminative materialism 36, 77, 117, 197, 293, 345, 436
- embodied 26, 78, 88, 96, 125, 186, 212, 266; action 29, 142, 205, 214; agents 311; cognition 77, 311, 314, 322, 323, 473; *see also* 4E cognition
- embodiment 455
- empathogens 351
- enactive cognition 70, 473
- encephalisation quotient 265
- encephalitis 150, 430
- endorphins 364, 430
- Engel, Andreas 133, 134, 136
- Engler, Jack 497
- enlightenment 175, 363, 490, 493, 497, 502–505, 506
- entheogens 355, 359
- enzyme 355, 359
- epigenesis 253–254
- epilepsy 133, 144, 382, 422, 440
- epileptic 379, 382, 424, 426
- epiphenomenalism 19–20, 43–44, 98, 197, 281, 284, 288
- epiphenomenon 25, 42, 114, 279, 370, 407
- episodic memory 150, 331, 353, 356, 400, 437
- epistemic 386, 437, 468
- epistemology 2, 481
- epoché 22, 472–473
- Erasmus 242
- Escher, Maurits Cornelis 294, 451, 452
- ESP *see* extrasensory perception
- eureka moment 213, 306
- evoked potential 79, 225–228, 425
- evolution: of consciousness 44, 207, 249, 276–284, 289, 298, 403; cultural 213, 255, 270; directed 251–254; of dreaming 403–404; theory of 19, 250, 278; zombie 41–44, 281–283, 287
- evolutionary psychology 276–278, 284
- evolutionary theory 249, 276, 335
- exceptional human experience 362
- existentialism 359, 374
- exorcism 122
- experimental psychology 25, 175, 447
- explanatory gap 28–29, 120, 183, 287, 504
- extended: consciousness 265, 449; mind 77, 125, 213
- extrasensory perception (ESP) 389
- eye movement 49, 130, 167, 406, 420

face: analysis of 89; area 91, 193, 384;
 recognition 146; visual 96
 false awakening 413
 false belief 267–269, 499, 505
 false intuition 49
 false memory 374–376
 ‘fame in the brain’ 117, 167
 Faraday, Michael 236, 238–239
 Farthing, G. William 346
 feature integration theory 132
 Fechner, Gustav 22, 160
 FEFs *see* frontal eye fields
 Feinberg, Todd E. 107, 258, 260, 285, 287
 Fenwick, Peter 363, 429, 506
 Feynman, Richard 354
 filling in the gaps 57–63, 449
 first-person approaches 465–469, 476, 478,
 482, 487
 First World War 421
 Fisher, Matthew 121
 Flanagan, Owen 281, 404, 478
 flash: drag effect 153; -jump effect 153; -lag
 effect 153
 Flash Gordon 283
 Flaubert, Gustave 35, 329
 Fletcher, Pat 209, 210
 flow 212
 flying 412, 415
 fMRI 80, 88–89, 162–163, 232–233;
 scanner 426
 form constants 381–384, 393, 412
 Four Noble Truths 490–491
 fovea 58–59, 65–66, 160, 162, 251, 331, 395
 Fox, Oliver 417
 Francis, Bruce 150
 Frankish, Keith 294
 Franklin, Stan 331, 462
 free will 118–121, 170, 218–245; as illusion
 53, 235, 238, 241, 243–244; *see also* agency
 free won’t 231–232
 Freud, Sigmund 30–31, 401, 404
 Friston, Karl J. 88, 406, 407
 Frith, Chris D. 223, 295
 frogs 258–259, 425
 Frohlich, Friedrich W. 153
 frontal eye fields 162, 164, 166, 181
 frontal lobe 78, 81, 135, 138, 212, 404
 fronto-parietal network 86–87, 131,
 224, 451
 fugue 33, 440, 450
 functionalism 198–199, 206, 283–288,
 293, 476–478; computational 310; and
 identity theory 19; and self 459; virtual
 machine (VMF) 334
 functional MRI *see* fMRI
 future selves 459–461

 Gallagher, Shaun 455, 472
 Gallup, Gordon 264, 265, 314
 Gallup poll 314
 Galton, Francis 143
 gamma 132–133, 182, 364, 420, 476
 ganzfeld 384, 389–390, 392
 gap: explanatory 28, 29, 120, 183, 287,
 504; filling in the 57–63, 449; mysterious
 27–29; William James and 57
 Gardner, Allen 272
 Gardner, Beatrix 272
 Gazzaniga, Michael S. 144, 146, 147, 148,
 439, 450, 451
 genetic drift 251, 254
 Geraerts, Elke 375
 germ line 253–254
 getting out of bed 222
 ghost in the machine 16, 18, 292, 430, 449,
 457
 Gibson, James J. 70
 Gillett, Grant 442
 glial cells 79, 290
 Global Workspace Theory (GWT) 100,
 113–117, 131, 140, 169, 196, 331;
 neuronal 116
 God 7, 175, 180, 219, 244, 250, 279, 280, 321, 492
 Gödel 451
 Gödelian 325
 GOFAI *see* good old-fashioned AI
 Goodale, Melvyn A. 203, 204, 205, 208, 210,
 227; profile 203
 Good Friday Experiment 362, 370
 Goodman, Rod 314
 good old-fashioned AI 307, 310, 326, 335
 Google 222, 315, 317, 384, 385
 Goswami, Amit 120, 253
 Gould, Stephen Jay 255
 grand illusion 52–72, 91, 171, 273, 374
 Gray, Jeffrey 36, 45, 196, 288
 Graziano, Michael S. A. 167, 171, 290, 295,
 333, 453; profile 168
 great chain of being 252–253
 Green, Celia 414
 Greenfield, Susan 259, 284, 320
 Gregory, Richard 26, 55, 61, 231, 354
 Greyson, Bruce 428
 Gulf War syndrome *see* PTSD
 GWT *see* Global Workspace Theory
 gyrus: cingulate 78, 80, 162; frontal 86,
 222; fusiform 86, 89, 91, 144; inferior
 temporal 86; occipital 86

- habituation 181, 407
- Hacker, Peter M. S. 147–148
- Haggard, Patrick 221, 223, 233, 234, 240
- half-second delay in consciousness 224–228
- Hall, Calvin 400, 478
- Halligan, Peter 152, 293
- hallucinations 376–387; auditory 359, 379, 381–382; Launay–Slade Hallucination Scale 378; misperception 376; pseudo 376–377, 383; visual 354, 359, 361, 378–379, 381, 384
- Hameroff, Stuart 46, 121, 122, 333
- Hamilton, Sir William 159
- Harding, Douglas E. 499–500, 501
- hard problem: of consciousness 28–29, 44–50, 57, 70, 84–87, 100, 111, 141; insolubility of 46; non-existence of 117; pretty 48, 140–141; solving the 120, 287, 330, 474, 476, 509; tackling easy problems 46
- Hardy, Thomas 373
- harmaline 359
- harmine 359
- Harnad, Stevan R. 304
- Harré, Rom 442
- Harris, Samuel B. 238, 243, 245, 492, 493, 499, 500; profile 491
- Hart-Davis, Adam J. 391
- hasheesh (hashish) 21
- health: inessentialism 282; mental 180, 356, 365; physical 175, 180, 191, 253, 339
- Hearne, Keith 416, 419
- heautoscopy 421–422
- Hebb, Donald O. 84, 133
- heedlessness 499, 502
- Heidegger, Martin 471
- hemianopia 206
- hemifield neglect 151, 152
- heterophenomenology 471, 483–488
- Heyes, Cecilia M. 266, 269
- higher order 34, 43, 231; consciousness 258, 261, 286, 415; global states (HOG) 197; perception (HOP) 198; representation 386; syntactic thought (HOST) 197; thought (HOT) 198, 209, 257, 261, 331
- Hilgard, Ernest 365, 366
- Hinduism 175, 180, 186, 423, 430, 493, 503
- hippocampus 80, 144, 150, 288, 399, 406
- Hirai, Tomio 181
- Hobbes, Thomas 290
- Hobson, Allan 348, 349, 401, 403, 404, 406, 407, 408, 415; profile 402
- Hodgson, Richard 439
- Hodgson, Shadworth H. 200
- Hofstadter, Douglas R. 283, 290, 325, 451, 452, 453, 468, 500, 501, 503
- Holland, Owen 313, 314, 324, 333; profile 332
- holographic reality 429
- Homer 286
- Homo habilis* 290
- Homo sapiens* 291, 395, 458
- homunculus 56, 100, 110, 117, 214, 313, 449, 450; somatosensory 99; unconscious 84, 111
- Honorton, Charles H. 392
- HOP *see* higher order, perception
- Horikawa, Tomoyasu 410
- HOT *see* higher order, thought
- Houston, Jean 360
- hua tou* 180, 500–501; *see also* koan
- Hubbard, Barbara Marx 146, 253
- Hui, Neng 499, 501, 504
- Hume, David 103, 439, 443, 448, 454; profile 438
- Humphrey, Nicholas K. 72, 100, 141, 267, 276, 280, 283, 285, 289–290, 291, 292, 293–294, 295, 404; profile 288
- Humphreys, Christmas 490, 496
- Hurley, Susan L. 141, 199, 205
- Husserl, Edmund G. A. 22–23, 471, 472–473, 475, 503
- Hut, Piet 472, 473
- Huxley, Aldous L. 357
- Huxley, Sir Julian S. 252
- Huxley, Thomas Henry 19, 200, 252
- Hyman, Ray 391
- hypnagogic 365, 377, 382, 412, 413; imagery 377, 382, 411, 412
- hypnopompic imagery 411
- hypnosis 25, 152, 343, 375, 391, 422, 440; as ASC 344, 351, 365–366, 369–370
- hypnotisability 212, 392, 422
- hypnotise 177, 192, 365, 366, 369, 423, 440–441
- hypothalamus 78, 80, 150
- IAM *see* internal agent model
- iCat robot 338
- IDA *see* intelligent distribution agent
- idealism 20
- identity 19, 33, 103, 118, 227, 274, 454, 475; behavioural 42; disorder 441; personal 346, 365, 437, 438, 446; theorists 77, 97, 346; theory 19, 117
- ideo-motor action 218

- Ikkyu 180
Iliad 285–286
 illusion 40; body-swap 263, 481; café wall 54–55, 376; of free will 243, 300; of impossibility 67; *maya* 493; rubber-hand 263, 354, 426, 480; self as 449, 451–452, 461; unity as 142–143; visual 22, 53–54, 204, 245, 294, 376; *see also* grand illusion
 illusionism 50, 243–244, 283, 294, 295
 imagery: hypnagogic 365, 377, 382, 411–413; hypnopompic 411
 imagination 365, 372–374, 378–379, 386–390, 396–397, 407; *see also* reality discrimination
 imitation 59, 265, 269–271, 285, 298, 336–338
 imitation game 316–318
 immortality 274, 436, 439, 459, 508
 inattentional blindness 54, 65–68, 71, 163, 169, 174, 183, 259
 incubation 212
 ineffability 37, 362
 inessentialism 42, 44, 209, 281–282, 284, 287, 319, 470
 inhibition 224, 297, 356; reciprocal 90
 insula 86, 382
 Integrated information Theory (IIT) 20, 91, 93, 119, 138–142, 167, 296, 333, 361
 intelligence 45, 92, 140, 250, 261, 265, 273, 281; artificial 2, 14, 26, 198, 273, 305, 316, 330; phenomenal 339; social 290–291; without representation 267, 312–314
 intelligent distribution agent (IDA) 331, 332
 intentionality 23, 26, 325, 325–329, 337, 472, 484; real 308, 326, 328, 485, 507
 intentional stance 267, 336–337, 484, 486
 interactionism: dualist 18, 120, 129, 195, 231, 437; symbolic 291
 Internal Agent Model (IAM) 332
 Internal World Model (IWM) 332
 internet meme 298, 299, 373
 interpreter 146, 439, 450–451
 introspection 23–24, 286, 291–292, 309, 464–467, 484–487
 intuition 34, 47–49, 56, 114, 172; pump 36, 39
 Islam 175, 219, 244, 437, 492
 IWM *see* Internal World Model
 Jackendoff, Ray 85, 478
 Jackson, Frank C. 38, 40
 Jackson, J. Hughlings 380
 James, Henry 21–22
 James, William 20–21, 23–25, 27–28, 30, 36, 44–45, 55, 57, 85, 107, 128, 142, 159–160, 168–170, 187–188, 193, 218, 220–221, 238, 284, 292, 299, 344, 352–353, 362, 439–440, 442, 446–449, 454, 464, 475, 506; profile 22
 Jarvik, Murray E. 383, 384
 Jastrow, Joseph 188
 Jaynes, Julian 285, 286
 jhanas 363–364, 493
 Johnson, Samuel 26, 222, 478
 Joyce, James 299, 309
 Joycean machine 299
 Judaism 175, 180, 492
 Jung, Carl G. 25, 382, 401, 424
 KA *see* Kernel Architecture
 Kabat-Zinn, Jon 176, 498
 Kafka, Franz 7
 Kammann, Richard 389
 Kant, Immanuel 136, 447
 Kanwisher, Nancy 91, 95, 205
karma 349, 506
 KASPAR 338
 Kasparov, Gary 315
kensho 502–503
 Kepler, Johannes 56
 Kernel Architecture 330
 ketamine 93, 353–354, 423
 Kihlstrom, John F. 192
 Kim, Jaegwon 199, 200
 Kirlian photography 396
 Kirsch, Irving 366, 369, 370
 Kismet ('sociable robot') 323–324, 337–339
 Klüver, Heinrich 382
 koan 180, 475, 500–501, 504
 Koch, Christof 46, 58, 59, 60, 77, 81, 83, 84, 85, 86, 87, 88, 93, 110, 111, 118, 122, 132, 133, 136, 145, 169, 170, 171, 297, 458, 476, 478; profile 95
 Korsakoff's syndrome 149–150
 cosmic consciousness 213, 253
 Kosslyn, Stephen M. 108
 Kozuch, Benjamin 88
 Kramnik, Vladimir 315
 Kubla Khan 212
 Kummer, Hans 267
 Kurzweil, Ray 320, 325, 459
 LaBerge, Stephen 416, 418, 419
 Lady Lovelace's objection 323
 Lamarck, Jean-Baptiste 251, 252
 Lamarckism 252
 language: animal 82, 257, 264, 272–273; body 67, 269, 337; instinct 271; natural 331, 335

- lateral geniculate nucleus (LGN) 38, 86, 90, 130, 154, 206, 453
 Latto, Richard 232
 laughing gas *see* nitrous oxide
 Launay–Slade Hallucination Scale 378
 Laureys, Steven 86, 348, 349
 Lavie, Nillie 165
 Lawrence, David Herbert 221
 Leary, Timothy 358, 361
 Lem, Stanislaw H. 321
 Leonardo da Vinci 56
 Leucippus 303
 Levine, Joseph 28, 468
 levitation 378
 Lewis, Clarence Irving 36
 Libet, Benjamin 18, 152, 169, 224, 225, 228, 229, 230, 233, 234, 235, 237, 479
 Libet’s Theories: backwards referral 226–227, 232; conscious mental field (CMF) 18, 45, 129, 196; half-second delay 224, 226, 228; theory of neuronal adequacy 225, 227–228; time-on theory 227; voluntary action 229–231
 LIDA 331–332
 life force 284, 293
 limbic system 80, 144, 430, 449
 limen 189; *see also* subliminal
 Llinas, Rodolfo R. 132
 Lloyd, Dan 77, 126, 474
 lobster 257, 261–262
 Locke, John 447
 locked-in syndrome 92, 191, 460
 Lodge, David J. 38–39
 Loftus, Elizabeth F. 375
 Logothetis, Nikos 82
 Lovelace, Ada 323
 LSD 179, 358–362, 368, 376–377, 382–384, 423; bicycle ride 360
 lucid dream 348, 398, 405, 408, 414–422, 430; inducing 416
 lucid dream machines 419
 Lucretius (Titus Lucretius Carus) 303
 Luna, Luis Eduardo 394
 Luria, Aleksandr Romanovich 143, 144, 150
 Lutz, Antoine 181, 474
 Lyell, Sir Charles 250
 lysergic acid diethylamide *see* LSD

 Mach, Ernst W. J. W. 139
 Machiavelli, Niccolò 267
 Machiavellian Hypothesis 267
 Machiavellian intelligence 271
 machine consciousness (MC) 303, 318–320, 324–325, 328–332; *see also* artificial consciousness
 machine modelling of consciousness (MMC) 308, 332
 machines: speaking 334–336
 Mack, Arien 65, 66, 163, 169, 174
 MacKay, Donald M. 147, 439, 449, 451
 Macphail, Euan 98, 257, 281
 macular degeneration 381
 magic 67, 68, 135, 220, 230; difference 31, 57, 84, 86, 91, 111, 187, 207, 288; ingredient 330; mushrooms 358
 magician(s) 67–68, 87, 236–237
 Maharishi Mahesh Yogi 175
 Malinowski, Peter 182
 Mallatt, Jon 88, 258, 260, 285, 287
 Maloney, J. Christopher 39
 Malthus, Thomas Robert 250
Man a Machine (L’homme machine) 19
 Mandler, George 240
 mantras 180, 503
 MAO *see* monoamine oxidase
 MAOs *see* monoamine oxidase
 mapping states of consciousness 347–351
 mapping the brain 79–80;
 electroencephalogram (EEG) 79, 133, 192, 348, 399–400, 405, 418–422, 429;
 nuclear magnetic resonance (MRI) 80, 89, 109; positron emission tomography (PET) 79, 93–94, 97; single cell recording 79, 86, 166, 223; transcranial magnetic stimulation (TMS) 80, 233, 350, 424–425;
 X-ray computed tomography (CT) 79;
see also brain, scanning
 Marcel, Anthony J. 190
 marijuana *see* cannabis
 Marino, Lori 265
 Mark 19 robot 40
 Marks, David F. 389
 Mary the colour scientist 38–41, 45, 282, 328, 470
 Masters, Robert E. L. 360
 matching-content doctrine 88
 materialism 19–20, 38–39, 48–49, 227, 304, 438; eliminative 36, 77, 117, 197, 293, 345–346, 436
 Matus, Juan 394
 Maury, L. F. Alfred 408, 411
maya 493
 MBSR *see* mindfulness-based stress reduction
 MC *see* machine consciousness
 McCarthy, John 328–329
 McEwan, Ian R. 80
 McGinn, Colin 1, 29, 45, 327, 330, 478
 McGurk effect 137, 386
 McKenna, Terrance K. 359, 383

MCS *see* minimally conscious state
 MDMA *see* ecstasy
 Mead, George Herbert 291
 mechanical Turk 305–306, 315–316
 medial lemniscus 225–227
 meditation 22, 175–184, 350, 362–364,
 494; altered states 351, 362–364;
 attention 175–184; basic principles
 177–180; Buddhist 495, 503;
 concentrative 177, 178–181; deep 198,
 300, 348; mindfulness 368, 419, 466, 495;
 nondirective 183; open or receptive 178,
 179, 181; posture 176; relaxation 175,
 176, 347, 349; transcendental (TM) 175,
 362, 503
 mediums 152, 236, 377–378, 423
 medulla 78, 79
 Meijer, Peter 209, 210
 Meltzoff, Andrew N. 269
 memeplex 271, 298–300, 461
 memes 213, 297–301; animal 270–271,
 273; creativity and 213, 298, 324;
 definition 271, 298; internet 298, 299,
 373; machine 300, 334–336; memeplex
 271, 298–299, 300, 461–462; and minds
 297–301; religious 219, 255, 299; self
 and 299–300, 334, 458, 462; selfplex
 299, 300, 334, 336, 462; tremes 299;
 viral 299
 memetic evolution 271, 335
 memetics 298–300, 335
 memory: and attention 166;
 autobiographical 114, 265, 375, 450;
 episodic 353; false 374–375; loss
 149–150, 155, 347, 354; man (*see also*
 Luria, Aleksandr Romanovich); semantic
 353, 356; short-term 115, 133–134, 166,
 309, 346, 356, 454; trans-saccadic 63;
 very short-term (VTSM) 173; working
 113–115, 131, 166–169, 187, 224,
 353, 420
 mental: field 196; function 347, 351, 363,
 365, 408; health 5, 180, 356; illness 2,
 14, 25, 243, 343, 351, 365–369, 379,
 396, 441, 462; model 307, 313, 456;
 science 21, 25; screen 108–111, 214, 288;
 states 19–20, 29, 43, 77, 187, 197–198,
 266–268, 437; theatre 104–106, 113,
 115, 456
 Mercier, Charles A. 27
 mere exposure effect 192
 mereological fallacy 77, 88, 148, 215, 235
 Merikle, Philip M. 190, 191, 193
 Merker, Bjorn 260
 Merleau-Ponty, Maurice 25, 26, 69, 70, 98,
 311, 471
 mescaline 357–358, 382, 384, 394, 422
 mesmerism 440
 Metzinger, Thomas 36, 129, 135, 136, 152,
 271, 290, 320, 332, 350, 408, 423, 426,
 452–453, 468, 504; profile 350
 Metzler, Jacqueline 108, 109
 Metzner, Ralph 361
 micro-awakenings 399, 416
 microconsciousness 136
 microtubules 46, 121–122, 324, 333
 midbrain 78–80, 92, 137–138
 Mikulas, William L. 494, 497
 MILD (mnemonic induction of lucid
 dreaming) technique 419
 Miles, James 243–244
 Miller, George A. 25
 Miller, Stephen M. 47–48
 Milner, A. David 203, 204, 205, 208,
 210, 227
 mind-body 199; connection 19; dualism
 98, 130, 294, 436; problem 11, 29, 32,
 100, 294, 388
 mindfulness 176–178, 182–183, 368, 498,
 505–507; *see also* meditation
 mindfulness-based, cognitive therapy 368;
 stress reduction (MBSR) 176, 498
 mindless design 249–251
 mind reading 236–237
 mindsight *see* change, detection
 mind-stuff 28
 mind-wandering 181–183
 minimally conscious state 86, 93, 405
 Minsky, Marvin L. 313
 mirror self-recognition (MSR) 263, 266–267
 Mitchison, Graeme 404
 Mithen, Steven J. 290–291, 292
 MMC *see* machine modelling of
 consciousness
 mnemonic induction of lucid dreaming
 416
 mnemonist 143
 modernist literature 21
 Mondrian 135
 monism 438; anomalous 15; dual-
 aspect 479; neutral 20, 28; reflexive
 477–479
 monist theories of consciousness 15, 18,
 21, 230, 353, 508
 monitoring: inner activity 151; open 177;
 reality 374, 379; self 43, 282, 295;
 source 375
 monoamine oxidase (MAO) 359
 Monroe, Marilyn 58–59

- Moody, Raymond A. 428
 Moody, Todd C. 42
 Moore's Law on Integrated Circuits 308
 moral: decisions 194–195, 278–280;
 responsibility 219, 232, 303, 437
 morality 7, 30, 231, 244, 278–280, 304
 Morland, Antony B. 207
 morphine 500
 Morse, Melvyn L. 424
 motor cortex 79, 80, 107, 230, 233, 237,
 399; primary 78, 222
 Moutoussis, Konstantinos 193
 movement 61, 67, 100, 131–132, 163, 222;
 eye 49, 69–71, 86–87, 109, 162, 166–167;
 involuntary 39, 80; rapid eye (REM)
 399–407, 409–410, 413–414, 418–420;
 saccadic eye 63, 162; skilled 201–202;
 voluntary 163, 228–235, 238–239
 movie-in-the-brain 55, 450
 Movshon, J. Anthony 85
 MPD *see* multiple, personality disorder
 MRI *see* fMRI
 MSR *see* mirror self-recognition
 Müller-Lyer illusion 54, 376
 multiple: drafts theory 90, 122–126, 142,
 155, 369, 386, 409–410; personality
 disorder 135, 143, 439–443, 450, 458
 realizability 47, 48, 198, 306; selves 458
 multisensory: dreams 406; integration 129,
 136–138, 429; perception 137, 429
 mushrooms, magic 355, 358, 394
 Myers, Frederic W. H. 378, 442
 myoclonic jerk 413
 mysticism 29, 44–45, 122, 318–320,
 476, 478
 mystical experiences 358, 362, 366, 424,
 494, 503, 508
- Nagel, Thomas 1, 32, 33, 34, 44, 257, 461,
 468, 477, 478
 narcolepsy 382, 413–415
 narrative gravity 336, 456–459
 natural selection 250–256, 276, 280–283,
 287–297
 nature's psychologists 293
 NCCs *see* neural correlates, of
 consciousness
 NDEs *see* near-death experiences
 near-death experiences (NDEs) 382, 387, 398,
 420, 423, 427–429; interpreting 429–430
 Near-Death Experience Scale 428
 Necker cube 81–82, 280
 neglect 58–59, 129, 143, 163; hemifield
 151–152; unilateral 151–152
- neocortex 80
 neo-Darwinism 254
 neural: Activation Mapping Project 236;
 Darwinism 296–297; networks 26,
 309–310, 324, 384–385
 neural correlates (NCs): of attention 166; of
 awareness 349; of consciousness (NCCs)
 xiii, 47, 81–88, 93, 174, 193, 260, 324,
 402; of experience 474; of free-will 224;
 of pain 96–97, 99; of vision 57, 205
 neurocentrism 77
 neuroeconomics 223
 neuromodulator 11, 348, 351, 399
 neuronal adequacy 225–228
 neuronal GWT 87, 116
 neurophenomenology 92, 181, 473–477,
 480, 495
 neurosis 395
 neurotransmitter 11, 308, 322, 348–351,
 364, 403, 406; dopamine 322, 423;
 glutamate 93; serotonin 322, 348
 niche construction 213
nirvana 349, 491, 506
 nitrous oxide 93, 352–353
 NMDA *see* N-methyl-D-aspartate
 N-methyl-D-aspartate (NMDA) 93, 353
 nociceptive signals 48, 99, 450
 Noë, Alva 17, 54, 56, 69, 70, 71, 72, 77, 88,
 91, 100, 101, 205, 313, 385, 456
 non-causal theories of consciousness
 197–198
 nonduality 180, 478–479, 493–495, 508
 noradrenaline 348, 353, 364, 399
 no-self 244, 436, 439, 493, 495, 498,
 505–509; *see also* anatta
 NovaDreamer 419
 nuclear magnetic resonance (MRI) 80, 109
 numbsense 206
- Oakley, David A. 293, 294
 OBEs *see* out-of-body experiences
 objective: reduction 120, 122, 324;
 threshold 190
 occipital lobe 78, 81, 86, 162
 O'Hara, Kieron 48
 Olson, Jay 236–237
 Omega Point 252
 ontogeny 48
 operant conditioning 24
 opium 21
 Orch OR (orchestrated objective reduction)
 120
 O'Regan, J. Kevin xiii, 56, 64, 65, 69, 70, 71,
 72, 100, 110, 313, 385

Origin of Species 250
 Ornstein, Robert E. 313
 Orwellian revision 112, 153–155
 ouija board 238
 out-of-body experiences xii–xiii, 152, 354, 359, 380, 413, 416, 420–427; asomatic 422; parasomatic 421; theories of OBEs 423–424; virtual reality 426–427; *see also* near-death experiences
 Own Body Transformation Task 425

 Pahnke, Walter N. 362
 pain 46, 96–101, 128, 196, 292; in animals 98, 262; in meditation 362–364; phantom limb 98–99; relief 352
 Paley, Reverend William 249, 250, 251
 panadaptationism 255
 panpsychism 20, 46, 136, 284
 Papineau, David 48, 49, 118
 paranoia 356, 392
 paranormal 25, 211, 236, 378, 387, 389–390, 431, 508
 parapsychology 378, 387–388, 392, 416
 parasomatic OBEs 421
 Parfit, Derek 437, 443, 506
 parietal lobe 78, 81, 86–87, 162, 416, 420
 Parkinson's disease 150, 423
 Parnia, Sam 428, 429
 Pascal, Blaise 305, 316
 Pashler, Harold (Hal) 159
 pathetic fallacy 49
 Pavlov, Ivan P. 24
 PCC *see* posterior cingulate cortex
 P-consciousness *see* phenomenal consciousness
 pearls on a string 453–454
 Peirce, Charles S. 36, 188
 Penfield, Wilder G. 424
 Pennartz, Cyriel M. 118–119
 Penrose, Sir Roger 46, 120–121, 122, 227, 294, 310, 324, 333
 Perky, M. Cheves W. 374
 Persinger, Michael A. 424
 persistent vegetative state (PVS) 86, 93
 PET *see* positron emission tomography
 PGO waves *see* pontine-geniculate-occipital waves
 phantom limb 98–100
 phenomenal: self-model (PSM) 320, 332, 408, 452; stance 337
 phenomenal consciousness 134–136, 198, 200, 294, 462, 469; in ASC 367; comparison with access consciousness 34–35, 88; definition of 34–35; as illusion 294–295; split 148; suffering 261; *see also* access consciousness
 phenomenality *see* phenomenal consciousness
 phenomenology 22, 46, 50, 124–125, 333, 471–479
 phenotechnology 350–351
 philosopher's zombie 41–44
 phi phenomenon 153
 phlogiston 293
 photosynthesis 276, 466
 phylogenetic 258–259
 phylogeny 48, 258–259
 physicalism 38–41, 198, 471
 pictorialism 56, 108–110, 385
 pigeon 24, 244, 265
 Pigliucci 47
 pineal gland 15–17, 56, 85, 105–107, 195
 Pinker, Steven A. 45, 78, 146, 255, 272, 277, 278, 335, 468
 Pino, Samanta 122
piti 363
 pixie dust 122
 PK *see* psychokinesis
 Plato 103, 186, 303, 427, 436, 447
 Plutarch 445
 pons 78–79, 399
 pontifical neuron 85, 107, 135
 pontine-geniculate-occipital waves (PGO) 406–407
 Ponzi schemes 299
 Ponzo illusion 376
 pop out 59–60, 71, 143, 146, 163, 167
 Popper, Karl R. 18, 120, 129, 195, 196, 227, 437
 positron emission tomography (PET) 79
 posterior cingulate cortex (PCC) 182, 401, 475
 postsynaptic membrane 309
 post-traumatic stress disorder (PTSD) 162, 352, 368, 392
 Povinelli, Daniel J. 263, 266, 268, 269
 precognition 387–388
 precognitive carousel 237
 precuneus 87, 181, 183, 356, 420
 prediction machine 26, 313
 pre-lucid dreams 414; *see also* lucid dream
 Premack, David 267, 269
 premotor 169, 221, 233; cortex 222; theory of attention 164, 166, 167
 preSMA 95, 222, 224
 pretty-hard problem (PHP) 48, 140–141

- primary: motor cortex 78, 222; visual cortex, (V1) 38, 61, 78, 86, 135, 162–163, 207, 361, 383
 Prince, Morton H. 440–442, 450
Principles of Psychology 21
 Prinz, Jesse J. 45, 85, 122, 318, 320
 problem-solving 144, 280, 284–285, 293, 356, 399, 419
 Probo 338
 procedural learning 149, 331
 propositionalism 108–109
 prosthetic limbs 322, 339
 protoconsciousness 404, 406
 proto-self 331, 449
 pseudo: -hallucination 376–377, 383; -randomness 310
 psi 388–390
 psilocybin 358, 361–362, 368
 PSM *see* phenomenal, self-model
 psychedelics 351, 354–361, 368, 385, 422, 508
 psychic: phenomena 378, 387, 440; powers 379, 389
 psychical research 377, 378, 387, 416, 422–425, 427, 440
 psychoactive drugs 200, 350–360, 368
 psychoanalysis 30, 186, 401, 404, 424
 psychodynamic theory 30
 psycho-integrator plants 394
 psychokinesis (PK) 388
 psycholitics *see* psychedelics
 psychonauts 359, 383–384
 psychoneural identity 227
 psychons 195
 psychophysics 18, 22, 95, 392
 psychotherapy 200, 317, 401, 496–498
 psychotomimetics *see* psychedelics
 PTSD *see* post-traumatic stress disorder
 pure consciousness 23, 179, 472, 503–504
 Puthoff, Harold E. 388
 PVS *see* persistent vegetative state
 Pylyshyn, Zenon Walter 110
 pyramidal cells 80–81
- QRio 319, 323
 qualia (singular 'quale') 35–37, 38, 41–47, 60–63, 77, 85, 91, 117; definition of 36; ineffable 37; of pain 98, 468; visual 206–207
 quantum: coherence 46, 120–122, 333; computers 121, 325, 333, 459; interactive dualism 120; phenomena 122; physics 46, 120, 122; theory 46, 121, 227, 253, 333
 quining qualia 36
- radical empiricism 25
 Radin, Dean I. 387
 Rahula, Walpola 506
 Ramachandran, Vilayanur S. 58, 59–61, 63, 98, 99, 100, 144, 146, 205, 449; profile 60
 Randi, James 67
 random number generator 390, 426
 rapid eye movement 399, 405, 416
 rat 24, 137, 254–256, 405, 406, 430
 rationality 168, 305, 315
 readiness potential (RP) 79, 221, 229–230, 233
 reality discrimination 374–376
 reciprocal altruism 279
 reduction, phenomenological 22, 472–473
 reductionism 28, 83, 97, 122, 277, 476–501
 reductionist theories 118, 294
 reflexive monism 477–479
 Reid, Thomas 438–439
 reincarnation 175, 439, 491, 496, 506
 Reiss, Diana 264
 relativity 37
 relaxation 175, 344, 349, 356, 364, 422; alert 176; deep 175, 346–347
 religion 7, 219, 252, 298–299, 422–423, 491–496; memes 219
 REM *see* rapid eye movement
 Rensink, Ronald A. 64, 65, 69, 70
 representational theories of consciousness 197–198, 313, 503
 repressed memory 375
 reptilian brain 80
res cogitans 304
res extensa 16
 retention, selective 251, 299
 retina 22, 38, 56, 61–62, 111, 130–131, 174, 234; blind-spot on 58, 62, 373
 retinal implants 208, 460
 retroactive masking 226
 retrograde amnesia 149
 retro-selection theory of dreams 410, 415
 reverberant loops 93
 Revonsuo, Antti 366, 367, 368, 402, 403, 404; profile 402
 Rhine, Joseph Banks 25, 387–388
 Rhine, Louisa 387–388
 Ricourte, George A. 352
 Rilke, Rainer Maria 56
 Ring, Kenneth 429
 rivalry 81–82, 88–90; binocular 62, 82, 90–91, 386
 Rizzolatti, Giacomo 166
 RNG *see* random number generator

RoboDennett 39, 40
 robot: anthropomimetic 322, 324,
 332–333; ethics 320; humanoid
 323–324, 338–339; lumbering 255;
 social 323, 337–339, 430; wall-following
 312, 314, 337
 robotics: behaviour-based 312, 327, 332;
 situated 312, 322; swarm 312, 322;
 theory-led 333
 Rock, Irvin 66, 163, 169, 174
 Romanes, George J. 269
 Rosch, Eleanor 496
 Rose, David 117
 Rosenberg, Monica O. 166
 Rousseau, Jean-Jacques 455
 RP *see* readiness potential
 rubber hand illusion 263, 354, 426, 480
 rule-and-symbol AI 307, 310
 Ryle, Gilbert 16–17, 18, 457

saccade 59, 63–65, 68–70, 162, 166, 167;
 sound 209; visual 405; voluntary 201
 Sacks, Oliver 149, 150
samsara 491, 496, 506
 Sand, Nicholas 359
 Sargent, Carl 392–393
 Sartre, Jean-Paul C. A. 59, 471
 satori experience 502
 Saunders, Guy 33
 Schank, Roger C. 326
 Schenk, Thomas 205
 Schilbach, Leonhard 481
 schizophrenia 162, 237, 353, 366, 379–381,
 386, 422
 schizotypy 422
 Schooler, Jonathan W. 243
 Schrödinger equation 120
 scotomas 61
 Scutt, Tom 48
 séance 236, 377–378
 Searle, John R. 244, 259, 271, 308, 315, 318,
 322–323, 325–329, 335, 468, 473, 478;
 profile 326
 second-person: methods 2, 465–467,
 472, 481–483; neuroscience 481–483;
 science 6, 124, 175
 second replicator 300
 Second World War 94, 253, 306–307, 499
 seizures 144, 349, 365, 382, 440
 Sejnowski, Terrance J. 153
 selective serotonin reuptake inhibitors
 (SSRIs) 351, 353
 self 17, 18; awareness 119, 181, 257, 265,
 415, 443, 449, 475; conscious mind
 129, 195–196, 227, 437; continuity
 of 449–451, 454, 508; egos, bundles
 and multiple selves 439; esteem 191;
 forgetfulness 180; and free will 29–30,
 504–505; monitoring 43, 282, 295;
 report 88, 182, 357, 502; sense of 6, 295,
 330–331, 420–421, 425, 455; supervisory
 system 147, 439, 451; theories of 439,
 443, 445, 455
 self-conscious 25, 129, 195–196, 227, 262,
 291, 356
 self-consciousness 7, 26, 135, 152, 263,
 271, 324
 self-control 223, 346, 440
 self-harm 365
 self-help 191
Selfish Gene 255
 selfish gene theory 255
 selfish replicator 254–256
 self-model theory of subjectivity 453
 self-monitoring zombies 43, 282, 295
 selfplex 299, 300, 334, 336, 462
 self-recognition 262–266, 267
 self-supervisory system 147, 439, 451
 self-transcendence 498
 selves in cyberspace 461–462
 senile dementia 150, 379
 sensorimotor: contingencies 62, 70–71,
 313, 385; theory 69–72, 109, 173, 294,
 313, 385
 sensory: cortex 78, 81, 86, 99, 107, 210,
 226, 232; deprivation 365, 380, 383,
 386, 408; isolation tanks 350, 354, 380;
 modality 53, 119, 137, 143; substitution
 138, 208–210
 sentience line 286
 serotonin 322, 348, 351, 353, 359–361,
 399, 405
 Seth, Anil K. 50, 260
 Seventh Sally 321
 Sewell, Anna 256
 Seyfarth, Robert M. 267
 Shakespeare, William 186, 219, 319
 shaman 359, 382, 393, 394
 Shannon, Claude E. 306
 Sheldrake, Rupert A. 390
 Shelley, Mary W. 305
 Shepard, Roger N. 108, 109
shikantaza 178
 short-term memory 115, 133, 134, 166,
 309, 346, 356, 454
 Sidis, Boris 188
 Siegel, Ronald K. 359, 380, 382, 383, 384,
 393, 414; profile 380

- signal detection theory 188–189
- simulation: disembodied 324; theory 266, 403
- Singer, Wolf 133, 135, 136
- single cell recording 79, 82, 86, 166, 223
- situated robotics 312, 322
- Skinner, Burrhus F. 24
- Slade, Peter D. 377, 378–379
- sleep: borders of 377, 411–412, 414; deprivation 365, 380; non-REM 399–403; onset REM 413; paralysis (SP) 397, 399, 411–414, 421, 430; REM 348–349, 399–410, 416, 419–420, 423; slow-wave 399
- sleep-wake cycle 28, 273
- sleep-walking 348, 365, 399, 418
- Sligte, Ilja G. 173
- Sloman, Aaron 258, 330, 334
- SMA *see* supplementary motor area
- Smithies, Declan 168
- snakes 152, 256, 377, 385, 394
- SoC *see* states of consciousness
- social: constructionism 26, 455; intelligence 267, 290–291
- Society for Psychical Research (SPR) 377–378
- society of mind 313
- sociobiology 277
- somatosensory cortex 78, 81, 86, 96, 99, 225–226, 232
- Soon, Chun Siong 233–234
- soul 7, 17, 19, 21, 44, 85, 186, 196; animal 425; and astral body 423–424; *atman* 493, 505; God-given 279, 280, 321–322; immaterial 17, 235; immortal 436, 439; and OBE/NDE 423–424, 429; the quantum soul 120; theory 447
- soundscapes 209–210, 419
- source amnesia 94
- SP *see* sleep, paralysis
- spatial: ability 114, 143, 278, 424, 429; attention 166–167; layout 69; map 137, 167; resolution 79, 160, 174
- speaking machines 334–336
- specious present 475
- Spencer, Herbert 28
- Spering, Miriam 87
- Sperling, George 172, 173, 174
- Sperry, Roger W. 144, 145, 146, 147, 230
- spinal cord 78–79, 96, 98–99, 128, 221–225
- spinocerebellar tract 79
- spirit 236, 239, 321–322, 377–378, 393, 414; animal 17, 200; reincarnating 437, 439
- spiritual experience 25, 360, 362, 494
- spiritualism 236, 378
- spirituality 2, 25, 428–429, 492, 498, 502, 504
- split brain 129, 135, 143–149, 443, 458; twins 149
- split personality 440
- spontaneous: awakening 494, 499; OBEs 426; trait transference 193; voluntary act (Libet) 229–231
- spotlight: of attention 113–114, 159, 161–167, 172; of focal consciousness 115
- Sri Aurobindo 252
- SSRIs *see* selective serotonin reuptake inhibitors
- SSS *see* state-specific sciences
- Stalinesque revision 112, 153–155
- Stapp, Henry 120
- Stargate project 389
- states of consciousness: altered (ASCs) 343–371; meditation and 175–184, 350, 362–364, 494
- state-specific sciences (SSS) 353
- Stazicker, James 172, 173
- Steels, Luc 335–336
- Stevenson, Robert Louis 440
- stimulants 351–352
- strange loops 451–452, 456, 501
- Strawson, Galen 453, 454, 501
- stream of: consciousness 21–22, 145–146, 444–448, 505; multiple drafts 123; vision 6, 54–57, 68, 71–72, 142
- stress: MBSR 176, 498; PTSD 162, 352; reduction 175–176, 183–184, 498
- string of pearls 453–454
- Stroop: effect 182, 190; interference 143
- Stuart, Susan A. J. 330
- subconscious 30, 333, 441
- subcortical: area 86, 138; loop 348; pathway 225; structure 208
- subjective: awareness 108, 202, 231; experience 37–41, 113–118, 121–122, 225, 229, 319; ontology 468; pain 96–97; qualia 35–36, 207, 288; referral in space 228, 402; referral in time 152–154, 226–228, 232
- subliminal: advertising 191; perception 31, 187, 188–189, 214; self 188, 191; stimulation 191–192; visual processing 86
- subsumption architecture 312, 334, 339
- suffering: animal 261–262; human 491; machine 324

Sufism 175
sukha 363
 super: -activation 295; -alertness 295;
 -blindsight 208–211
 superego 30
 superior colliculus 86, 130, 137–138, 154,
 162–164
 supernatural 45, 220, 235, 253, 357,
 379, 495
 superunity 143
 supplementary motor area (SMA)
 221–222
 supraliminal perception 189
 survival: after death 7, 219, 378, 430,
 437, 459, 508; and evolution 289, 292,
 297; of the fittest 251, 255; value of
 consciousness 287–289
sutras 491
 Swiss Army Knife 277–278, 290–291
 symbolic: filling-in 59, 62; species 291;
 thought 271, 291
 symbol manipulation 326, 328
 synaesthesia 60, 129, 137, 143–146, 209,
 288, 354–356, 360
 synaesthetes 143–146, 288, 382
 synapses 18, 45, 120–122, 296, 304,
 308–310
 synchronisation 338
 synchronised oscillations 46, 133–134
 synchrony 101, 232, 475; binding by
 132–135, 171
 syndrome: amnesic 150; anarchic hand
 223; Charles Bonnet 381, 384, 386;
 chronic fatigue 367; Korsakoff's
 149–150; locked-in 92, 191, 460;
 philosophers' 40; Tourette's 224
 Szilard, Leo 213

 tactile vision substitution systems (TVSS)
 208
 Tai Chi 194, 350
 talking heads 335–336
 Tallon-Baudry, Catherine 133
 Tamagotchis 336
 Taoism 497, 507
 Targ, Russell 388, 389
 Tart, Charles, T. 344, 353
 Taylor, John, L. 331
 Teilhard de Chardin, Pierre 252, 253
 telepathy 129, 378, 387–389
 teletransporter 444–445, 448, 451
 Teller, John, T. 67
 temporal lobe 61, 78, 107, 162, 358, 413,
 424, 430
 temporoparietal junction (TPJ) 183, 361,
 413, 425
 terminally ill 358, 360, 362
 termites 261, 270, 312
 tetrahydrocannabinol (THC) 355–356,
 382–384
 tetrahydroharmine 359
 thalamocortical: loops 80, 93, 118, 138;
 system 93, 132, 260, 297, 399, 407
 thalamus 78–80, 86, 93, 150, 225, 358, 450
 THC *see* tetrahydrocannabinol
 Theatre *see* Cartesian Theatre
 theories of consciousness: causal and
 non-causal theories 195–198; enactive
 theories 141, 293; functionalism
 198–199, 285, 288, 293, 476; Global
 Workspace Theory 113–115, 167, 169,
 196–197; integrated information 20,
 93, 119, 138–141; materialism 19–20,
 197, 214; multiple drafts 90, 122–126,
 386; mysterianism 45, 320, 476–478;
 predictive processing 294, 386, 406;
 reductionist 83–84, 118, 294, 476;
 representational theories 198–199, 367;
 stream of vision 6, 54–57, 68–72, 142;
 unity in action 141–142; visual binding
 46, 131, 133; without theatres 117–122
 theory of mind (TOM) 266–269, 273, 285,
 481; computational 308; interaction 266,
 481; simulation 266, 403
 Theosophy religion 423
 therapy 360, 441, 496–499; behavioural
 367, 498; cognitive 368, 498; depression
 358; psychotherapy xv, 200, 401, 496
 Theravada 363, 491
 thiamine 150
 third-person 173, 474, 478, 482–483; data
 466–468, 476, 504; methods 465–466,
 477–478; observations 477, 481;
 perspective 97, 200, 287, 417; science
 175, 465; views 4–6
 third replicator 299–300
 thisness 13
 Thompson, Evan 88, 101, 108, 475, 492,
 509
 Thorndike, Edward Lee 270
 thought: insertion 236; suppression 177
 thought experiments 44, 304, 322, 459,
 477; blindsight 209; Chinese Room 308,
 325–328; intuition pumps 36; Marilyn
 Monroe 58–59; Mary the colour scientist
 38–41, 45; rules of 42–43; sensorimotor
 theory 70, 72; what is it like to be 257;
 with self 443–446

- time 226; backwards in 226, 232; displacement 153; half-second delay 224–226
- time-on theory of consciousness 227
- timing and volition 232–234; *see also* attention
- Titchener, Edward B. 23, 204–205, 465
- Titchener illusion 204–205
- TM *see* transcendental meditation
- TMS *see* transcranial magnetic stimulation
- Tolstoy, Count Lev Nikolayevich 412
- tomography: positron emission (PET) 79; X-ray computed (CT) 79
- Tononi, Giulio 93, 119, 138–139, 296, 297, 333, 404; profile 138
- Tooby, John 277
- total therapy 497
- Tourette’s syndrome 224
- Tower of Generate-and-Test 296–297
- TPJ *see* temporoparietal junction
- trance 146, 152, 365–366, 378, 440, 498
- transcendence 25, 180, 253, 362, 459, 498; of self 498
- transcendental meditation 175–176, 179–180, 348, 362–363, 503
- transcranial magnetic stimulation (TMS) 80
- trans-saccadic memory 63
- Treisman, Anne 132
- trimethoxyphenylethylamine 358; *see also* mescaline
- Triplett, Norman 67
- Troscianko, Emily T. xii, 123; profile 4
- Troscianko, Jolyon T. 261
- Trurl’s perfection 321–322
- tunnels and lights 430
- Turing, Alan M. 306, 307, 310, 314, 319, 321–323, 324, 325; profile 307
- Turing machine 307–310, 314, 321, 323
- Turing test 307, 314–319, 325–328, 333
- Turk, mechanical 305–306, 315–316
- TVSS *see* tactile vision substitution systems
- Tyndall, J. 27
- unconscious: action 193, 194–197, 203–205, 220; emotional and social effects 192–193; inference 26, 30, 187; muscular action 236, 238; perception 188–194, 214; plagiarism 94; previewing 90; processing 90–91, 105, 148, 186–188, 210, 227–228
- unconsciousness 92–96, 192, 205, 423, 462
- unilateral neglect 151–152
- unity: in action 141–142; of consciousness 119–121, 128–130, 130–155; as illusion 130, 142; of self 132, 435, 446–449, 455; superunity and disunity 143; in time 152–155
- Universal Darwinism 251, 255, 296–297, 298
- Universal Turing Machine 307
- user illusion 300, 334
- Utts, Jessica 389
- van de Castle, Robert L. 400
- van Eeden, Frederik W. 415
- van Lommel, Pim 387, 428–429, 430
- Varela, F. 91, 473–474, 475, 476–477, 478, 495; profile 473
- Vedanta 497
- vegetative state 85–86, 93, 348–349
- Velmans, Max 97, 200, 287, 477–479, 485, 493
- velocity of thought 22
- ventral: allocentric system 184; pathway 89, 205; stream 81, 91, 131, 162, 203–205, 208, 227, 259
- ventriloquism 137
- ventro-lateral prefrontal cortex 181
- ventromedial prefrontal cortex (mPFC) 182, 223
- very short-term memory (VSTM) 173
- veto, conscious 231
- Vicary, James 191
- viral memes 299
- virtual machine functionalism (VMF) 334
- virtual reality 99, 392, 407, 415, 453; generator 404, 406–407; out-of-the-body 426–427
- virtual representation 69–70
- virus: of the mind 299; religion as 299
- vision: conscious 57, 83–84, 86, 142, 207, 209; stream of vision 6, 54–57, 68–72, 142
- visual: areas 62, 133, 144; binding 46, 131, 133; illusion 22, 53–54, 204, 245, 294, 376; perception 53, 56, 69, 160, 203–204, 386; stimuli 87, 173, 204, 225
- visual cortex 33, 82, 86, 107, 109, 111, 130; activity in 13, 163, 210; in cats 132; damage to 135, 162, 379; primary (V1) 38, 61, 78, 86, 162, 207; tunnels and lights 430
- visualisation 180, 495
- visuomotor control 132, 194, 201, 203–205, 208
- vitalism 322
- VMF (virtual machine functionalism) 334
- Vohs, Kathleen D. 243

- volition: neuroanatomy of 222–224; and timing 232
 voluntary action 221–224, 228–231, 233–234, 506
 von Helmholtz, Hermann L. F. 22, 26, 30, 164, 165, 187
 von Leibniz, Gottfried W. 304, 305
 Voorhees, Burt 122
 VSTM *see* very short-term memory
 Vygotsky, Lev S. 26, 291

 Wagstaff, Graham 366
 Walter, W. Grey 237, 313, 362
 Warrington, Elizabeth K. 206
 Watson, John B. 24, 25
 Watt, James 213
 Watts, Alan W. 244, 497, 506
 Watzl, Sebastian 168, 172
 Webb, Taylor W. 167, 168, 171
 Weber, Ernst H. 22
 Weber–Fechner Law 22
 Wegner, Daniel M. 83, 95, 177, 188, 232, 238–240, 243, 244, 507; profile 238
 Weir Mitchell, Silas 98
 Weiskrantz, Lawrence 150, 151, 206, 209
 Weismann, August F. L. 253
 Wei Wu Wei 244
 Wernicke’s area 78, 107
 West, Donald J. 378
 West, Louis 380
 what it is like: to be a bat? 32, 44, 257, 461; to be an octopus? 257, 273; to be something? 27, 29; to be you? 4, 35
 Wigner, Eugene P. 120
 Wilber, Kenneth E. 252–253, 497
 Wilberforce, Samuel 252
 Williams, George C. 255, 277
 Williams, Mark 498
 Williams, Venus 202
 Wilson, David Sloan 255
 Wilson, Edward O. 255, 277
 Wiseman, Richard J. 378, 419
 Wittgenstien, Ludwig J. J. 16, 289
 Woodruff, Guy 267, 269
 Woolf, A. Virginia 5
 Worsley, Alan 416
 Wren-Lewis, John 500–501, 503
 Wundt, Wilhelm M. 23, 24, 153, 165, 465

 ‘X’ 330–333
 X-ray computed tomography 79

 yage *see* ayahuasca
 Yamdoots 430
 Yānomamö 393
 yoga 176, 180, 194, 348, 350, 363, 419, 497
 yogin 347, 497, 504

 Zahavi, Dan 454, 455, 471, 475
 Zeki, Semir 107, 135, 136, 148, 154, 193
 Zen 178, 244, 353, 362–364, 495–503;
 and attention 180–181; and the brain 181; Buddhism 180, 363, 493, 495–496; enlightenment 502–504; paradox 495, 501, 504; and perplexity 501; training 181, 497
 Zener cards 388
 zimbo 43, 282–283, 295, 457
 zombic hunch 285, 467, 469
 zombie: evolution of 281–283, 287; impossibility of 41–42; partial 206–207; philosopher’s 41–44; replica 41; self-monitoring 43, 282, 295; twins 41–42, 282



Taylor & Francis Group
an informa business

Taylor & Francis eBooks

www.taylorfrancis.com

A single destination for eBooks from Taylor & Francis with increased functionality and an improved user experience to meet the needs of our customers.

90,000+ eBooks of award-winning academic content in Humanities, Social Science, Science, Technology, Engineering, and Medical written by a global network of editors and authors.

TAYLOR & FRANCIS EBOOKS OFFERS:

A streamlined experience for our library customers

A single point of discovery for all of our eBook content

Improved search and discovery of content at both book and chapter level

REQUEST A FREE TRIAL

support@taylorfrancis.com

 **Routledge**
Taylor & Francis Group

 **CRC Press**
Taylor & Francis Group