Attention

Wayne Wu



New Problems of Philosophy

ATTENTION

Attention is a fundamental feature of the mind yet has languished in the backwaters of philosophy. Recent years, however, have witnessed a resurgence of philosophical interest in attention, driven by recognition that it is closely connected to consciousness, perception, agency, thought, justification, and introspection. As is becoming clear, attention has a rich philosophical significance.

This is the first book to provide a systematic overview and assessment of different empirical and philosophical aspects of attention. Wayne Wu discusses the following central topics and problems:

- major experiments and theories of attention in psychology since the 1950s
- the neuroscience of attention, including basic mechanisms and models
- attention's intimate relation to agency
- the phenomenology of attention
- attention as a gatekeeper for consciousness
- attention as the basis for perception-based thought about objects
- the role of attention in the justification of belief
- attention in introspection of consciousness.

A key feature of the book is its skilful analysis of the empirical work on attention, and how this relates to philosophy. Additional features include chapter summaries, annotated further reading, and a glossary, making this an ideal starting point for anyone studying attention for the first time, as well as being suitable for more advanced students and researchers in psychology, cognitive science, and philosophy.

Wayne Wu is Associate Professor and Associate Director of the Center for the Neural Basis of Cognition, Carnegie Mellon University, USA. He was previously Assistant Professor at The Ohio State University, USA.

NEW PROBLEMS OF PHILOSOPHY

Series Editor: José Luis Bermúdez

"Routledge's New Problems of Philosophy series has a most impressive line-up of topical volumes aimed at upper-level undergraduate and graduate students in philosophy and at others with interests in cutting edge philosophical work. The authors are influential figures in their respective fields and notably adept at synthesizing and explaining intricate topics fairly and comprehensively." – John Heil, Monash University, Australia, and Washington University, St Louis, USA

"This is an outstanding collection of volumes. The topics are well chosen and the authors are outstanding. They will be fine texts in a wide range of courses." – Stephen Stich, Rutgers University, USA

The New Problems of Philosophy series provides accessible and engaging surveys of the most important problems in contemporary philosophy. Each book examines either a topic or theme that has emerged on the philosophical landscape in recent years, or a longstanding problem refreshed in light of recent work in philosophy and related disciplines. Clearly explaining the nature of the problem at hand and assessing attempts to answer it, books in the series are excellent starting-points for undergraduate and graduate students wishing to study a single topic in depth. They will also be essential reading for professional philosophers. Additional features include chapter summaries, further reading and a glossary of technical terms.

Also available:

Analyticity Cory Juhl and Eric Loomis

Fiction and Fictionalism *Mark Sainsbury*

Physicalism Daniel Stoljar

Noncognitivism in Ethics Mark Schroeder

Moral Epistemology Aaron Zimmerman **Embodied Cognition** *Lawrence Shapiro*

Self-Knowledge Brie Gertler

Semantic Externalism Jesper Kallestrup

Consequentialism Julia Driver

Images John V. Kulvicki Forthcoming:

Imagination Fabian Dorsch

Emergence Patrick McGivern

Egalitarianism Iwao Hirose

Disjunctivism Matthew Soteriou

Cognitive Phenomenology *Elijah Chudnoff* **Perception** Adam Pautz

Metaphysics of Identity André Gallois

Modal Epistemology Otávio Bueno and Scott Shalkowski

Social Metaphysics Amie L. Thomasson This page intentionally left blank

ATTENTION

Wayne Wu



First published 2014 by Routledge 2 Park Square, Milton Park, Abingdon, Oxon, OX14 4RN

And published by Routledge 711 Third Ave., New York, NY. 10017

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

© 2014 Wayne Wu

The right of Wayne Wu to be identified as the author of this work has been asserted 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

Library of Congress Cataloging in Publication Data Wu, Wayne. Attention / by Wayne Wu. -- 1 [edition]. pages cm. -- (New problems of philosophy) Includes bibliographical references and index. 1. Attention. 2. Consciousness. 3. Philosophy of mind. I. Title. B808.9.W8 2014 128'.3-dc23 2013043302

ISBN: 978-0-415-53222-8 (hbk) ISBN: 978-0-415-53224-2 (pbk) ISBN: 978-1-315-77760-3 (ebk)

Typeset in Joanna by Taylor & Francis Books

CONTENTS

	List of Figures	viii	
	Acknowledgements	x	
Introduction: The puzzle and challenge of attention			
1	The Psychology of Attention	11	
2	The Neuroscience of Attention	45	
3	Agency and the Metaphysics of Attention	76	
4	Attention and Phenomenology	110	
5	Attention as the Gatekeeper for Consciousness: Inattentional Blindness	146	
6	Attention as the Gatekeeper for Consciousness: Cognitive Access	176	
7	Attention and Demonstrative Thought	209	
8	The Epistemic Role of Attention	243	
Conclusion: What attention is and why it is central			
	Appendix: Shannon Information Theory Glossary Bibliography	275 280 287	
	Index	309	

LIST OF FIGURES

1.1	Switching in Dichotic Listening	18
1.2	Visual Search	22
1.3	Feature Integration Theory	24
1.4	Spatial Cueing Paradigm	25
1.5	Reaction Time Benefits of Cueing	26
1.6	Time Course of Effects in Direct and Indirect	
	Spatial Cueing	27
1.7	Two Attentional Networks	28
1.8	Yarbus and Eye Tracking	35
2.1	Marrian Explanatory Framework	49
2.2	Plots of Neural Activity (Spikes)	51
2.3	Receptive Fields of IT and V1 Neurons	52
2.4a	The Human Visual System	53
2.4b	The Macaque Monkey Visual System	53
2.5	Multiplicative Gain	55
2.6	Contrast Gain	56
3.1	A Behavior Space and the Many-Many Problem	80
3.2	Where is Attention as Selection for Action?	96
4.1	Contrast and Spatial Frequencies	116
4.2	Gabor Patches	117
4.3	Carrasco Contrast Experiment	118

4.4	Carrasco Landolt Square Experiment	120
4.5	Tse Illusion	121
4.6	Nine Squares	133
4.7	Two Gray Dots	134
5.1	The Gatekeeper and Common Sense Models	147
5.2	Crowded Lines	168
5.3	Identity Crowding	169
6.1	The Gatekeeper and Common Sense Models	176
6.2	Cognitive Access and the Gatekeeper and Common Sense Models	177
6.3	The Flow of Information via Working Memory	179
6.4	The Neural Global Workspace	182
6.5	Letter Array in Sperling Partial Report Paradigm	188
6.6	Contrast and Spatial Frequency	196
7.1	Box-line figure used by J. Duncan to probe object attention	212
7.2	Object-based Attention Paradigm	213
7.3	Multiple Object Tracking (MOT) Paradigm	214
7.4	Ishihara Diagrams	220
8.1	The Direct Model of Introspection	256
8.2	The Transparency Model of Introspection	258

ACKNOWLEDGEMENTS

Since I began this project, I have imagined an ideal book on attention that I might write. This book is not that book, but it is closer to it because of many colleagues and friends who have helped me in numerous ways.

A number of scientists whom I have never met in person received spontaneous queries from me and were kind enough to respond, often in detail. These include Charles Folk, Liqiang Huang, Victor Lamme, Sid Kouider, Hal Pashler, Anne Treisman, and Jeremy Wolfe. Bob Kentridge also shared his knowledge of many relevant issues and provided helpful feedback as a referee. Alan Allport, who I believe got many of the main issues about attention right, discussed matters with me some years ago via email. My own work on the nature of attention develops Allport's seminal contributions.

I'm very grateful for interactions over the past few years with Chris Mole and Declan Smithies. I have learned much from their work on attention. As colleagues at The Ohio State University, Declan and I organized a conference on the philosophical significance of attention, and many of the papers presented at that venue were subsequently published in a volume Chris, Declan, and I edited. I learned much from all the participants at that conference, and I still regret that we did not have the space to publish all the articles presented there.

John Searle and Jay Wallace guided my own work on this topic as sympathetic and critical dissertation advisors. I'm grateful for their guidance and encouragement. John Campbell influenced me through his writing and in conversations with him in person and by email. Although I was never officially his student, John's instruction, philosophical example, and support have been greatly appreciated.

In 2012, I helped organize the Attention discussion group at the Center for the Neural Basis of Cognition (CNBC) at Carnegie Mellon. The group brought together students, post-doctoral researchers, and a large number of faculty members who are experts on many areas of psychology and neuroscience. I thank all the participants, but I want to highlight my faculty colleagues: Aaron Batista, Marlene Behrmann, Steve Chase, Marlene Cohen, Carol Colby, Raj Ghandi, Rob Kass, Tai Sing Lee, Matt Smith and Byron Yu. A web search on their work should make clear how lucky I was to have access to their expertise. Several graduate students and post-docs also aided my thinking in this context including Erin Crowder, Adam Greenberg, Marvin Leathers, and Linda Moya. I thank the other participants who contributed to a lively, productive discussion on low-level to high-level aspects of attention. They showed me what ideal cognitive science is. I also want to single out Carl Olson and Marlene Cohen for patiently helping me with divisive normalization (and of course, any errors are mine alone).

Many people read parts of the book, gave me comments on ideas presented in the book, and/or hashed things out with me in extended conversations, as well as by email: Marlene Behrmann, Ned Block, Felipe de Brigard, Ben Bronner, John Campbell, Imogen Dickie, Steve Chase, Todd Ganson, Brie Gertler, Chris Hill, Rob Kass, Bob Kentridge, Philipp Koralus, Giulia Martina, Chris Mole, Bence Nanay, Adam Pautz, Ian Phillips, Sarah Shomstein, (who also had her lab read chapters!), Nico Silins, Declan Smithies, James Stazicker, Henry Taylor, John Tsotsos, and Sebastian Watzl. I want to especially thank Hemdat Lerman, who read through most of the book and gave me detailed comments on nearly every chapter. John Schwenkler invited me to post on my work at the Brains blog, and I'm grateful to him and to the participants who engaged with me on attention and introspection. Jake Beck had his cognitive science seminar go through several chapters as they discussed attention, and I am grateful for comments from his students and from Jake that have helped me improve things. Mark Wheeler stepped in to generate two images of the human brain in an hour of great need, for which I'm grateful. I also thank the graduate participants of my attention course: Ori Beck, Stephen Foley, Shivam Patel, and Laura Tomlinson. Later, Brian Ballard, Ori, and Laura worked carefully through many of the chapters at the final stage, helping me with last minute improvements and corrections as well as raising a host of challenges that talented graduate students are wont to do. I'm grateful for their attention. I'm also grateful to Cameron Wittig (http://cameron-photo.com/) for providing the very neat photos of the Ishihara figures in Chapter 7.

Tony Bruce first approached me with the idea of writing on attention for the New Problems of Philosophy Series. I'm grateful to his entrusting me with this task. Adam Johnson then guided me through the editorial process and answered my many niggling questions. Later, Andrew Watts aided me through the production process. I'm grateful for their responsiveness, aid and encouragement. Three anonymous referees also provided me with detailed comments that helped me make substantial improvements. I'm grateful for their efforts. Peter Murray's copy-editing morphed into helpful substantive editing, and I'm grateful for his queries which improved the manuscript at the final stage.

In many books, the most important acknowledgments always seem to come last, but for me, that is because I want my deepest gratitude to be what reverberates at the end. Alison Barth helped me with some of the basic neuroscience, and my daughters, Madeleine and Eleanor, diverted me and kept me going (Eleanor even drew one of my figures). All three provide encouragement and meaning simply by existing in my world. This book is dedicated to them, the attentional highlights of my life.

INTRODUCTION THE PUZZLE AND CHALLENGE OF ATTENTION

Attention insinuates itself into what we do, perceive, and think. When you listen carefully for the location of a beeping noise, examine the complex texture of an object with your fingers, try to isolate the hint of strawberry in a red wine, or just look for your keys, you deploy attention, sometimes successfully as when you see your keys, sometimes not, as when you fail to taste the hint of strawberry in the wine. Attention can also involve movement, as when you use your fingers to search for something under the couch or shift your eyes as you follow someone moving across a room, but it can also be furtive, as when you are pretending to be interested in a conversation yet secretly listen to the conversation behind you, or when you look at someone "out of the corner of your eye." Moreover, attention is not always voluntary. It can be involuntary, as when a loud noise or awful smell pulls your focus to it. Even in thought, attention plays a role: the embarrassing event of yesterday that constantly pulls your thoughts to it, distracting you throughout the day; the attempt to remember the name of a person and then, finally recalling it; or just thinking, working through specific reasons to make a decision and thereby following a train of thought. Attention is a fundamental capacity of the mind, and it is what this book is about.

Given the pervasiveness of attention in human behavior, it is not a surprise that it has been the subject of intense experimental investigation in empirical psychology and in neuroscience. Recently, philosophers have begun to focus more attention to attention. This is appropriate, for as this book argues, attention is of rich philosophical significance. Attention insinuates itself into a host of phenomena of central concern to philosophers: agency, thought, inference, introspection, perception, and consciousness. A few vignettes will provide anchor points:

Inattentive and attentive driving

Jane is coasting along in her car, deep in thought. She expertly shifts gears, changes lane, makes turns, all while thinking about something else. Suddenly, she notices that there is more traffic and shifts attention to driving, attending to obstacles and expertly weaving in and out between the cars. As she clears the clutter, she goes back to autopilot and switches attention back to the previous train of thought.

The hidden surprise

Jane is further up the hiking trail than Sam, and peers around the bend. "Sam, you've got to look at this!" she shouts. Sam, huffing and puffing well behind her, breaks into a run. "Hurry!" Jane calls, "You've got to see this!" "What is the thing that Jane sees?" Sam wonders. "Is it dangerous?" Meanwhile, Jane focuses on the glowing, undulating object, wondering: "What is that?"

The difficult decision

Should she quit her job? Jane focuses her energies on attentively considering a slew of reasons for and against. She finds her work boring. At the same time, the pay is excellent and she has the authority she long worked for. Her office, however, leaves much to be desired, located in a noisy wing without decent windows, though that seems a silly reason to contemplate leaving. There are so many reasons, but to make a good decision, Jane will have to base it on good reasons.

The beauty of fuschia

Sam has never seen fuschia. He has heard it described but guesses that it is perhaps a florid yellow. Then he sees it for the first time, carefully attending to it and marveling at its color. Even more so, he marvels at that qualitative aspect of his experience of fuschia: phenomenal fuschia. "I wish I had experiences more like that!" he exclaims, attentively introspecting the phenomenology of his experience.

Gates and spotlights

Jane is conversing with Sam, but her thoughts drift to her big decision. Should she quit and accept the offer from her company's main competitor? Should she negotiate for a better package with her current company? She has been invaluable on the last few major projects. But ... "Jane?" Sam's voice intrudes. Surprised, Jane realizes that Sam is still there. It was as if he had disappeared, lost in the haze of her thoughts, but suddenly, looking directly at him, Jane realizes that he had been talking all that time.

Each vignette provides an instance of a phenomenon that has triggered intense and comprehensive philosophical reflection: the nature of intentional, skilled, and reflexive agency; the role of perception in thought; the responsiveness to reasons in deliberation; the epistemic and subjective properties of introspection; the scope, structure, and limits of consciousness. In each case, attention permeates the phenomenon. This book provides an argument that attention is not just pervasive but plays a fundamental role in these central aspects of mind. Accordingly, philosophical accounts of these features will not be adequate unless they come to grips with attention.

To address the question of philosophical significance, we begin with a more basic question, the *metaphysical* question: What is attention? After all, how can we propose adequate theories that invoke attention as a central component, let alone ask clear questions about attention's role, without having a good idea of what it is?

Given the amount of work done in psychology on attention, we might look for an answer there. It is appropriate to begin with one of the founders of modern psychology, William James, who famously noted:

Everyone knows what attention is. 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, and is a condition which has a real opposite in the confused, dazed, scatterbrained state which in French is called distraction, and Zerstreutheit in German.

(James 1890, 403)

This passage is full of ideas, but at the outset, it seems unhelpful in providing us a theoretical account of attention. Rather, James voiced what he took to be common knowledge, what ordinary people take attention to be. His description, then, needs to be unpacked if it is to provide an answer to our metaphysical question.

We might hope that a century after James, we can answer the metaphysical question, but this is not so. Many contemporary theorists think dimly of the concept of attention, and this pessimism seems to be growing. The negative assessment of our conceptual understanding has, apparently, been long-standing. Christopher Mole cites Karl Groos (1898), who noted a century ago that

To the question, "What is Attention?", there is not only no generally recognized answer but the different attempts at a solution even diverge in the most disturbing manner.

(cited by Mole, 2010, p. 4)

Indeed, one suspects that many contemporary attention theorists will find little to disagree with when Jeremy Wolfe and Todd Horowitz claim that "there is no single, satisfying definition of attention" (Wolfe and Horowitz 2004). Stuart Sutherland, in a review of two monographs on attention, notes:

Over the past 50 years, the sheer ingenuity displayed by psychologists working on attention rivals if it does not exceed that of cosmologists studying black holes. Indeed, there is a similarity in their results – after many thousands of experiments, we know only marginally more about attention than about the interior of a black hole.

(Sutherland 1998, 350)

Felipe de Brigard and Jesse Prinz assert that

those who try to move beyond that suggestion that "everyone knows what attention is" often replace the folk concept with idiosyncratic definitions

that settle crucial questions by fiat rather than facilitating the process of scientific investigation and discovery.

(de Brigard and Prinz 2010, 52)

It will not be hard to find other psychologists and philosophers making similar remarks about the state of attention research. Some of my empirical colleagues have noted darkly that attention is what cognitive scientists invoke when they encounter something inexplicable in their data. For example, an unexpected result must be due to some shift in attention. The current state of play has prompted one theorist, Britt Anderson, to argue that there is no such thing as attention (Anderson 2011). This is unsettling. How can the theory of attention have strayed so far from James's certainty?

This book pursues a more optimistic outlook on answering the metaphysical question. From the extant discussions, there are three families of proposals to consider regarding what attention is:

- 1. A function-centered approach: what attention is for.
- 2. A mechanism-centered approach: how attention works or is implemented.
- 3. A phenomenology-centered approach: what attention is subjectively like.

In the pages that follow, we shall identify proposals in each of these families. On the function-centered approach, we shall examine the following, with emphasis on the visual modality:

- Attention as a filter of information for further processing
- Attention as binding features for object representation and awareness
- Attention as a spotlight (perhaps zoom-lens), highlighting its target
- Attention as selecting targets for memory, consciousness, or action.

In each of these, we have James' original idea of the selectivity of attention. On the mechanism-centered approach, the emphasis has been on the nature of the process of selection, described at different levels of analysis, from more abstract computational algorithms to concrete neural mechanisms. Among these, we shall consider:

- Attention as the modification of neural signals
- Attention as altering the area of space to which a neuron responds
- Attention as emerging from competition for limited resources
- Attention as the preparation of a motor response.

6 INTRODUCTION

We can treat these as claims about the realization of attention. Finally, on the phenomenology-centered approach, one follows William James' claim that consciousness is of the essence of attention and focuses on the phenomenology of attention echoed in talk of attention as like a spotlight. Thus, we consider

• Attention as a distinctive mode of consciousness.

The plethora of options is a bit dizzying. No wonder that theorists of attention doubt that there will be a single unified answer to the metaphysical question.

To different degrees, we will touch on each of the proposals just noted. It will be helpful, however, to examine the issues from the perspective of one of the answers to the metaphysical question. This book will settle on a conception of attention as the subject's selecting an item for the purpose of guiding action, what is called in the psychological literature the selection for action account of attention. This account focuses on attention's role in behavior. This is the conception of attention I have defended in other work. However, in preparing this book, and especially in thinking about experimental work on attention, I have been struck by how the seeds of the selection for action account are already found in basic assumptions within the experimental practice of psychologists and neuroscientists working on attention. Or so I shall argue. Accordingly, the selection for action account is neither idiosyncratic nor stipulative. Rather, it emerges from shared assumptions in cognitive science. The account strikes me as the best hope of imposing unity and organization on the theory of attention, and thus of answering the broad skepticism that we noted earlier. Still, philosophy being what it is, there are counterarguments and counterproposals. There are many open questions and much work to be done. It is my hope that this book will inspire readers to do further work.

It is important to recognize that while our aim is for a univocal account of attention, it might turn out that there are in fact too many disparate phenomena that we label "attention." Thus, the psychologist Alan Allport (1993), in a wide-ranging critical review of a quarter century of attention research up to the early 1990s, suggested that we begin

by taking seriously the idea that attentional functions are of very many different kinds, serving a great range of different computational purposes. There can be no simple theory of attention, any more than there can be a simple theory of thought. A humble but also a more ambitious task for the next twenty-five years will be to characterize, in cognitive neurobiological terms, as much as possible this great diversity of attentional functions. (207)

Perhaps he is right, but even if "attention" names a large family of phenomena, there is still a question about what unifies the members of each subfamily of attention, and for each, a different version of the metaphysical question can be formulated. So we cannot fully escape the question of metaphysics. In the space allotted, this book will pursue as much unification as can be squeezed out of philosophical argument informed by relevant empirical research. Allport's suggestion, however, should always be kept in mind.

A guiding principle of the book is that any theorizing about attention must come to grips with the central results of the cognitive science of attention. Accordingly, the first two chapters of the book spell out some of these results. This is to provide the reader both with a selective background in the vast, indeed overwhelming, empirical literature, but also with materials for fashioning an answer to the metaphysical question. It is important to note, however, that the metaphysical question has largely not been a concern in the empirical study of attention. Indeed, the question might seem of little importance to experimentalists, a way of obfuscating results obtained from experimental work. For their part, philosophers have sometimes ignored empirical work, perhaps seeing it as irrelevant to the abstract theoretical issues they investigate (this is changing).

It will be helpful to state five basic questions regarding attention, the first of which is the metaphysical question. The other four questions are as follows:

Function: What role does attention play?Properties: What are characteristic features of attention?Mechanism: How is attention implemented?Consciousness: What is the relation between attention and consciousness?

A complete theory of attention answers all five basic questions, and completing such a theory will require both empirical and philosophical work. But why should philosophers consider empirical work, and why should experimentalists consider conceptual work? Philosophical theories of the nature of mind have empirical implications. After all, if philosophers provide theories of action, of intention, of introspection, or of reasoning, those theories apply, presumably, to human minds as well. Yet human minds are also objects of empirical investigation, so to the extent that these philosophical theories say something about how the mind works, they are beholden to empirical work that speaks to that issue. If I am right that attention plays a central role in many of the mental phenomena of interest to philosophers, then philosophers theorizing about those phenomena have strong reason to heed the many empirical results regarding attention. What of the other direction? What do cognitive scientists have to gain from philosophical concerns such as the metaphysical question? It is true that the unavailability of an answer to the metaphysical question does not impede rigorous experimental investigation on attention. Even as cognitive scientists have bemoaned the lack of adequate definitions of attention, they have at the same time enriched our empirical understanding of how attention works, what it does, and how it might influence consciousness. A variety of well-defined experimental paradigms have yielded substantial knowledge of the properties and mechanisms of attention. Why then worry about metaphysics?

In providing an answer to the metaphysical question, by giving an account of what attention is, we fix the referent of the concept attention, and in doing so, anchor specific parts of our theories of the mind to the world. Without that anchor, theoretical invocations of attention will be like a boat adrift. While the experimental paradigms we shall discuss in later chapters have increased our understanding of attention, the use of these paradigms in experiments is necessarily tied to theory. For experiments are often conducted to answer specific questions in order to test hypotheses about attention, yet the adequacy and clarity of these questions and hypotheses depend on the clarity of the concept of attention. So, while the experimental methods might be clear, they are used in the service of answering theoretical questions. Cognitive science is not just in the business of probing nature by experiments, but also of constructing theories of psychological phenomena such as attention and other mental faculties. It is at the theoretical level that answering the metaphysical question matters, and it is here that philosophy and empirical cognitive science meet.

Let me now give a very brief summary of the structure of the book so readers can get started. Chapter 1 provides an overview of psychological work on attention since the 1950s, covering central questions, theories, models and experimental paradigms that have engaged psychologists studying attention. In particular, I shall identify a link between attention and tasks that will reverberate through the rest of the book. This chapter also provides some rigorous definitions of central concepts used to characterize attention. Chapter 2 then moves to the neuroscience of attention, focusing on different effects on neural activity when attention is deployed, on different models of neural processes, and on the possibility of drawing on mechanisms to answer our metaphysical question. I suggest, in the end, that the task-centered conception of attention uncovered in Chapter 1 provides a way of unifying disparate neural phenomena. Chapter 3 then takes this task-centered conception and develops it into a full-blown answer to our metaphysical question: attention is selection for action. By situating attention within agency, I hope to illuminate both phenomena. In Chapter 4, I turn to the phenomenal conception of attention, and consider whether consciousness is of the essence of attention. or at least, where attention is tied to consciousness, whether attention has a uniform phenomenology. While attention seems to affect phenomenology, it does not do so uniformly, and this raises questions for a phenomenal conception of attention. A different connection to consciousness is explored in Chapters 5 and 6: attention serving as a gatekeeper for consciousness. Specifically, I examine the claim that one is phenomenally conscious of X only if one attends to X. The experimental evidence, I shall argue, falls short of settling this matter, but more importantly, much conceptual work remains to be done to even understand what the claim is. In the last two chapters of the book, we will look at applications of attention in understanding demonstrative thought (Chapter 7) and in justification and introspection (Chapter 8). In all these chapters, I hope to underscore the philosophical importance of attention.

Much ground regarding attention will be covered, but there are many things that I have left out that I would like to have discussed in detail. Among them are: empirical phenomena of attention, such as inhibition of return and the attentional blink; disorders of attention such as hemispatial neglect, attention deficit hyperactivity disorder, and attention in mental disorders such as schizophrenia; attention and the binding problem; attention and ambiguous figures; attention in non-visual perceptual modalities; the many disparate computational approaches to attention; attention to time; attention and veridicality; attention and language use (pragmatics); attention and testimony; and joint attention (though on the latter, see the volumes by Eilan et al. 2005 and Seemann 2011). I also think the epistemic significance of attention is an area that will see rapid growth and the coverage of that issue is more cursory than I would like, in part because much exciting work remains to be done. Furthermore, the connection between attention and ethics is largely unexplored, though relevant ideas crop up in the discussion of virtue.¹ As it is, I hope that readers will find much in the book to engage them. Attention is a central capacity of mind, yet for so

long it has been neglected in philosophical theory. We are beginning to address this neglect, and we will do so in what follows. So, let's begin with a bit more attention to attention.

Note

1 On recent related work in ethics, namely, an account of modesty as a phenomenon of attention, see (Bommarito 2013).

1

THE PSYCHOLOGY OF ATTENTION

1.1 Introduction

This chapter considers highlights of psychological research on attention since the 1950s, a period during which the conceptual scheme that frames contemporary theorizing about attention was firmly established. It examines central experimental paradigms used to probe attention, the initial questions and theories that drove early investigation, some conceptions of what attention is, and the concepts developed to characterize attention. Aside from a historical overview, there are two additional goals. First, with an eye to answering the metaphysical question, "What is attention?", I shall extract from experimental paradigms a link between attention and a subject's selecting information or targets to guide and control performance of a task. Specifically, I argue that a background assumption in experiments on attention is that such selecting for task is sufficient for attention. This condition provides the seed of an answer to the metaphysical question to be developed in subsequent chapters. Second, these experimental paradigms have informed the development of a theoretical vocabulary to characterize attention, and in particular, have led to descriptions of two basic kinds of attention: roughly, (a) attention that can be intentionally directed as when one looks for a missing object; and (b) attention that is captured as when a loud sound pulls one's focus to it. I will provide a rigorous analysis of the concepts used to characterize this division.

12 THE PSYCHOLOGY OF ATTENTION

Section 1.2 begins with the common idea that attention is a form of selection but raises the question, "Selection for what?" Section 1.3 then examines an early debate about what stage of perceptual processing selection occurs at, in particular, whether it is at early or late stages of such processing. Here, attention was conceived of as a filter for information, selecting it for further processing. As the early versus late selection debate was never adequately resolved, section 1.4 discusses the proposal of Nilli Lavie's Load Theory of Attention that the conflicting data that drove the debate was a function of the different experimental tasks researchers used. The nature of the task makes a difference. In that vein, while early work focused on auditory attention, work on vision became prominent in the 1960s. Section 1.5 discusses the visual search paradigm and a resulting theory due to Anne Treisman: the Feature Integration Theory. While this theory is no longer at the center of current debates, it set the stage for how attention is conceptualized. Section 1.6 then discusses another paradigm for spatial attention, spatial cueing, and considers the contrast between top-down versus bottom-up attention. Are there different types of attention or different attentional mechanisms? Section 1.7 picks up on this theme and examines some central conceptual dichotomies used to characterize attention. I provide definitions of the central dichotomies. Finally, section 1.8 extracts from the standard experimental paradigms a sufficient condition for attention to an X: selection of X for a task. I argue that this is a shared assumption that can serve as an antidote to the widespread skepticism about an answer to the metaphysical question noted in the Introduction.

1.2 Attention as selection for what?

The Introduction presented five basic questions about attention:

Metaphysical: What is attention?Function: What role does attention play?Properties: What are characteristic features of attention?Mechanism: How is attention implemented?Consciousness: What is the relation between attention and consciousness?

To begin the discussion of the psychology of attention, consider the function question. There is widespread agreement among cognitive scientists that attention is a process of selection. James's passage captures the selectivity commonly attributed to attention: "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." Attention cannot, however, be merely selection. After all, there are many kinds of selection that do not count as attention. An object sorter can be highly selective yet does not attend to what it selects. As to be discussed in Chapter 2, a neuron can be highly selective in having a preferred stimulus, but it does not follow that the neuron thereby attends to its stimulus as opposed to its being part of a mechanism of attention. Indeed, there is something odd about the claim that a neuron, a part of a person, attends. The point is that if attention is selection, it is a specific kind. Psychologists often add that attention is selection for further processing, but this invites similar challenges: the object sorter and neuron can select for further processing, too. Further precision is needed in characterizing attentional selection.

One way to distinguish attentional selection from other forms of selection is to identify the type of thing that can attend. James speaks of the taking possession "by the mind" emphasizing that it is a psychological subject that pays attention, namely an entity that has a mind. The previously noted object sorter and selective neuron are not psychological subjects, so they cannot exemplify attentional selection even if they exhibit another kind of selection. One can then treat attention as a subject-level phenomenon or, as philosophers like to put it, a personal-level phenomenon. The relevant contrast is between the personal and the subpersonal. Although this distinction is widely invoked, it needs clarification. In the absence of a rigorous analysis of the distinction, I proceed with a simple division. One can think of personallevel states as those states that are attributable to a subject and not to the subject's parts, such as the brain or part of the brain. In contrast, subpersonal states are attributed to those parts but not to the subject. On this account, unconscious mental states count as personal in that they are attributed to the subject and not to the subject's parts. For example, certain parts of the brain might implement Freudian Oedipal desires, but while the subject might have such desires, that part of the brain does not. Similarly, attention is something that persons are capable of, not their parts. If one were to accept this division of the personal from the sub-personal, then one can discount selection exhibited by neurons and dumb machines as forms of attention.¹

Still, the answer is not very informative, for while it suggests what kind of thing can be selective, it does not tell us much about selection. Might there be something in the nature of attentional selection that also divides it from other kinds of selection? Let's begin the historical overview of the psychology of attention while holding this question constantly in the background.

1.3 The debate over early versus late selection: capacity limitations

In the revival of modern attention research in the mid-twentieth century, attention theorists focused on the selection of information: psychological subjects are presented with a lot of information in experimental situations, and, to perform a task, they must select only relevant information. For example, a subject asked to selectively listen to one of two conversations selects information from that conversation. This emphasis on information, inspired by communication theory in the 1950s, led to the first major debate about attention: At what point in perceptual information processing does attentional selection occur? The answers to this question, usually divided between so-called *early selection* and *late selection* accounts, provide an early account of what attentional selection involves.

It is common among cognitive scientists to speak of both the mind and brain as processing information, but what is information? Claude Shannon (1953) provided a precise definition in his theory of communication in terms of what he called mutual information. The latter is defined in terms of entropy, which in information theory is a statistical measure of uncertainty. This concept of information is defined mathematically, but I eschew the technical details and make do with three points: (a) mutual information is tied to the reduction of uncertainty (a message about X is informative to the extent that it reduces uncertainty about X); (b) information can be precisely quantified (often measured in bits, derived from "binary digits"); and (c) it is not identical to meaning: the same meaningful sentence can carry different amounts of mutual information, while two sentences of different meaning can carry the same amount of mutual information (for more on information, see Appendix A). Meaning can be understood as a type of semantic information where "semantic information" identifies the content of a representation. Such content need not be linguistic such as the content (meaning) of a sentence, but can also be tied to representations of features or objects, such as the auditory system representing pitch or the visual system representing a ball. Given the distinction between semantic and mutual information, an ambiguity crops up in talk of processing and, later, of selecting information. Does "information" mean mutual information or semantic information? In fact, in psychological and philosophical theorizing,

it is often the latter that is meant, but then what is the significance of mutual information?

In his book, Perception and Communication (1958), Donald Broadbent drew on Shannon's theory to propose a "fresh language" (35) and a "new set of descriptive terms" (36) for psychology. On Broadbent's view, the technical language of information allows for precise characterization of information processes that are capacity limited. These processes can only deal with a limited amount of information at a time. For example, Itti and Koch (2001) suggest that information can flow at $10^7 - 10^8$ bits per second along the optic nerve transmitting information from the eye. How can visual processing keep up with this vast input? Experience also suggests that there are capacity limits to perception. For example, there are a limited number of conversations you can listen to at once. It is natural to characterize this limit in terms of an informational bottleneck, although this is merely a metaphor. Given Shannon's work, Broadbent realized that psychology could go beyond metaphors to investigate capacity limits with mathematical precision. This is an important point that has been lost in recent years as psychologists and philosophers have focused on semantic information (meaning). Theorists have invoked capacity limits in theories of attention and, as we shall see, in theories of working memory in connection with phenomenal consciousness (see Chapter 6). They have suggested that such limits impose constraints on the nature of attention and consciousness, but in general, invocations of capacity often remain qualitative, rather than quantitative. These theories thus suffer the fate that Broadbent sought to avoid: metaphors rather than precision. Ultimately, serious talk of capacity limits must quantify these limits if the invocation is not to be merely a figure of speech. It will not be possible to invoke information theory in detail in our discussion. Rather, the point is to remember that where invocation of capacity limits becomes important, a theory must provide quantitative measures of the sort Broadbent drew from Shannon.

Capacity limits yield a plausible story of why attention is necessary. For, given a limited capacity to deal with an overabundance of information, a creature needs a capacity to select just what information is relevant for current goals on pain of information overload. That is, a capacity-limited creature needs attention. Capacity limits and selection provide the conceptual background for the debate over early versus late selection: Does attentional selection occur early or late in perceptual processing? The idea is that there are capacity limits on information processing, namely, the maximum amount of mutual information that can be processed at a time.

Plausibly, limits on processing of mutual information impose a limit on the amount of semantic information (representational content) that can be processed at a time. In what follows, we focus on limits in processing semantic information in light of a channel's limited capacity to process mutual information. Thus, we shall focus on the processing of representational contents, constrained by the (mutual) information capacity of the relevant processing channel. Talk of early and late selection concerns the different stages of perceptual processing of relevant representations. For example, in audition, an early stage of processing concerns the basic audible features of a sound (e.g., a voice)—say, its pitch or timbre—while a late stage of processing concerns the categorical features of the sound, say, the identity of the voice or the meaning it expresses. The question then is whether attention selects basic or categorical features.² Assuming that perceptual processing is capacity limited, an informational bottleneck must occur somewhere, and attention then serves to select information at the bottleneck. In light of this, Broadbent suggested that attention acts to filter information.

In early work on attention, pioneered by Colin Cherry (1953), the focus was on auditory processing of language. Cherry focused on filtering tasks where subjects are presented with multiple stimuli and asked to select some subset of them. In the dichotic listening paradigm, two streams of verbal inputs are presented, one to each ear treated as a separate information channel. Subjects then selectively "shadow," i.e., verbally repeat, only one of the sound streams. The basic finding was that when subjects attended to one stream, they did not pick up information from the other. When queried about what was said in the unattended stream, subjects were unable to provide accurate answers (notice that this experiment focuses on semantic information, i.e., what is heard, not mutual information). If perceptual processing was not capacity limited, psychologists initially reasoned, then subjects should be able to report the contents of both channels.

Jon Driver (2001) has noted that two questions were fundamental to theorists at that time: (a) What conditions allow people to effectively shadow the attended message?; and (b) What do people typically know about the unattended message? On the first question, it was ascertained that substantial differences in the physical properties of the sounds between the two channels facilitated performance. For example, shadowing improved when the two auditory streams were heard as if coming from distinct rather than the same locations. Shadowing was also aided by distinct acoustical properties such as presenting a low-pitch versus a high-pitch voice. In general, physical distinctness aided attentional selection.

Regarding the second question, many experiments suggested that subjects miss quite a lot from the unattended channel. Building on Cherry's work, Neville Moray (1959) observed that even when the unattended channel consisted of a small number of words repeated multiple times, subjects still failed to report accurately what those words were. In general, early observations suggested that while subjects could notice abrupt changes in lowerlevel physical properties of the unattended stream, higher-level perceptual properties like semantics (meaning) were typically missed. Consequently, Broadbent postulated that attentional filtering occurs after processing of basic physical features but prior to processing of categorical features. Thus, filtering occurs early in perceptual processing. The general picture entails a division between a preattentive and an attentive stage of processing, a distinction that remains to this day. On Broadbent's early selection account, preattentive processing concerns basic physical properties of the stimuli, with attention filtering relevant information about basic properties for higher-order, categorical processing. While talk of attention as a filter is metaphorical, Chapter 2 will consider one possible neural implementation of attentional filtering. For now, construe the metaphor as Broadbent's answer to the function question: attention filters (selects) information for the purpose of categorical processing. This provides a more concrete specification of attention as a type of selection.

Evidence quickly accumulated, however, showing that quite a bit of semantics in the unattended channel could get through. Drawing on the anecdote of the cocktail party effect where the mention of one's name in a nearby conversation is said to capture attention, Neville Moray (1959) did find that when the subject's name appeared in the unattended channel, the subject was more likely to notice it: subjects reported instructions expressed in the unattended channel when these were preceded by their name (33% of the time, Moray, op. cit. table IV, p. 58). Moray concluded that "[i]t is probably only material 'important' to the subject that will break through the [bottleneck] barrier" (op. cit., 56).

Subsequently, Anne Treisman (1960) argued that "the selective mechanism in attention acts on all [stimuli] not coming from one particular source by 'attenuating' rather than 'blocking' them" (246–7). In one experiment, Treisman instructed subjects to shadow a verbal stream presented to one ear, say the right ear, while ignoring a second stream presented to the other ear. Unbeknownst to her subjects, Treisman swapped the verbal streams midsentence, so a sentence that begins in the right ear, switches to the left, and vice versa. An example is given in Figure 1.1:



Correct Shadowing: "I saw the girl song was wishing"

Figure 1.1 Treisman's experiment where two verbal streams are presented to each ear (one in italics, the other in bold). Normally, each stream is presented to a single ear, but Treisman changed channels mid-sentence so that the verbal streams switched ears at the point indicated by the vertical line. Thus, the sentence, "I saw the girl jumping in the street," begins in the left ear, the attended channel, but jumps to the right ear. Subjects were asked to shadow only one of the channels, so correct performance is just shadowing of the words in a single channel. The arrows indicate what the subjects actually shadowed, and, speculatively, they suggest that the subject's attention jumps between channels despite task instructions.

Here the sentence to be shadowed jumps from the right ear to the left, with the switch point indicated by the vertical line. To shadow correctly, however, the subject must continue to repeat the words on the right. In the example, the two sentences at issue are as follows: first, "I saw the girl jumping in the street," which begins in the right channel but switches to the left after "girl"; and, second, "me that bird song was wishing," which begins in the left channel but switches to the right after "bird". Correct shadowing would be the nonsensical "I saw the girl song was wishing." Surprisingly, the subjects shadowed "I saw the girl jumping wishing". It was as if attention jumped between the two ears despite task instructions. Indeed, subjects were unaware of doing this. Treisman reasoned that this intrusion of semantics from the unattended channel on the shadowing of the attended channel depends on contextual effects that continue to influence behavior, since the word "jumping" rather than "song" is more probable given the preceding "I saw the girl ... " It seems that unattended information remains available to influence behavior and is not completely filtered out. Given Treisman's plausible explanation of why the subject jumps between

Actual Shadowing: "I saw the girl jumping wishing"

channels, it is natural to conclude that despite being unattended, the left channel must be analyzed to a higher linguistic level.³ This means that information selected by the filter is not the only information being processed at higher stages. Since this unattended information is not being fully blocked, one might wonder if the filter isn't leaky: unattended information can get through the filter for higher level processing. Alternatively, perhaps the filter is operating at a *late* stage in perceptual processing. Indeed, the latter possibility began to gain wider acceptance.

By the mid-1980s, Daniel Kahneman and Treisman (1984) noted a shift from early to late selection theories (Deutsch and Deutsch 1963; Norman 1968). On late selection accounts, filtering occurs after all signals are perceptually processed up to a categorical level of representation (e.g., semantic). Thus, relevant capacity limits occur post-perceptually, and it is only at this late stage that attention is needed. As Driver puts it

Late selectionists ... proposed that the limited awareness of unattended stimuli (as for the non-shadowed message in selective listening experiments) might have less to do with rejection from full perceptual processing, than with rejection from entry into memory or into the control of deliberate responses ... Thus, unattended stimuli might conceivably undergo full perceptual processing, yet without the person being able to base their deliberate responses upon this, and without the formation of explicit memories.

(2001, 58)

The point of late selection accounts is that perceptual processing may not be limited at all; rather the bottleneck occurs when perception engages other systems.

The debate about early versus late selection highlights some answers to the basic questions: attention is a type of selection, namely filtering; it occurs at specific moments in perceptual processing; and it functions to deal with capacity limitations. What is left hanging is whether attention operates early or late in perceptual processing.

1.4 Task demands and load: resolving early versus late selection

Kahneman and Treisman (1984) noted that the shift from early to late selection theories coincided with a shift in different types of experimental paradigms. Early work in audition involved filtering tasks where subjects were overloaded with task-irrelevant input. Later work focused less on filtering and more on target selection. This includes well-known paradigms like spatial cueing and visual search to be discussed in sections 1.5 and 1.6. Kahneman and Treisman argued that different experimental paradigms might tap into different mechanisms involving selection at different stages in processing. Thus, disparate results favoring early or late selection might merely reflect choice of experimental task.

Nilli Lavie and co-workers have proposed a Load Theory of Attention that builds on Kahneman and Treisman's observation (Lavie 2005). Begin with the idea of processing as resource limited, so that the amount of processing available for any task has an upper bound. Unlimited processing capacity is not available. What happens to the limited resource if current processing does not use all of it? Does the remainder lie dormant? Does another process tap into it? Treisman (1969) suggested that "we tend to use our perceptual capacity to the full on whatever sense data reach the receptors" (p. 296).⁴ In line with this, Lavie and Tsal (1994) suggested that total processing resources are always deployed, and where the attended information channel does not exhaust available resources, remaining capacity is then apportioned to processing unattended channels. On their account, what is critical is the perceptual load of the attended channel, namely how much of available processing resources that channel consumes. This suggests an explanation of the conflicting data that drove the early versus late selection debate. Load Theory holds that both models are in a sense correct, for the observed effects adduced to support either early or late selection depend on task demands, namely what the subject is doing. The general prediction is that early selection effects will be seen in high perceptual load conditions where all available processing is consumed by heavy task demands. For example, auditory filtering tasks in early work in the 1950s might be high-load, involving a large amount of information to be sifted through. In contrast, late selection effects will be seen in low perceptual load conditions where the system is not overloaded with information and additional processing resources are available for processing unattended channels.

The crucial point is that the nature of the experimental task can effect how attention is deployed, which can give rise to either early or late selection effects. The character of these effects is task dependent because tasks determine the informational load that must be processed. Thus, a potential resolution of the early versus late selection debate is that both are in a sense correct, with their characteristic effects differentially occurring depending on the perceptual load in the task. More importantly, if early and late selection effects depend on the nature of the task, then it would be incorrect to tie the selectivity of attention down to a specific stage in processing (i.e., early or late). Rather, a more general possibility is beginning to emerge: sometimes, attention can be early in processing; sometimes it can be late. In either case, attention is dependent on the task.

1.5 Visual search and the Feature Integration Theory of attention

While early attention research focused on audition and verbal shadowing tasks, there was a gradual shift to vision and visual search tasks in the 1960s. Visual search is looking for something. Sometimes, search is difficult, as when you look for a friend in a crowded train station; sometimes it is easy, as when that friend is wearing a neon green shirt, jumping up and down in plain view. The attention that guides visual search can be understood as directed at objects and/or their features.

An influential model of visual search was Treisman's Feature Integration Theory of visual attention (FIT) (Treisman and Gelade 1980; for an informative assessment, see Quinlan 2003). In its initial version, FIT treats visual object recognition as a constructive process where basic visual features are first detected by dedicated receptors, e.g., those for color, shape, and motion. The visual system then binds these features to form representations of objects. As in Broadbent's filtering conception of attention, visual object recognition involves two stages, in this case a preattentive feature detection stage and an attentional binding stage. Treisman and Gelade (1980) wrote that in FIT:

features are registered early, automatically, and in parallel across the visual field, while objects are identified separately and only at a later stage, which requires focused attention. The model assumes that the visual scene is initially coded along a number of separable dimensions, such as color, orientation, spatial frequency, brightness, direction of movement. In order to recombine these separate representations and to ensure the correct synthesis of features for each object in a complex display, stimulus locations are processed serially with focal attention. Any features which are present in the same central "fixation" of attention are combined to form a single object. Thus focal attention provides the "glue" which integrates the initially separable features into unitary objects. Once they have been correctly registered, the compound objects continue to be perceived and stored as such.

(98)

Accordingly, in the preattentive stage, processing of features occurs in parallel and independent of focused attention. Focused attention then binds features into object representations. A critical aspect of the model is that the processing of features and objects is separate.

Standard visual search tasks require subjects to search for a target amid a set of distractor objects, the number of distractors constituting the set size. There are two experimental conditions, target present and target absent. The subject reports whether a target was present or absent (yes, in the first condition; no, in the second). Furthermore, there are two kinds of search: feature search, where a single basic feature is the target, and conjunction search, where targets are individuated by a combination of basic features.⁵ For example, in feature search, where color and shape are basic features, one might look for a red **T** against a sea of green and brown **T**s (i.e., red is the relevant feature). In this case, one looks for a difference in color to identify the target. In conjunction search combining both shape and color, one might look for a green **T** in a sea of brown **T**s and green **X**s. In this case, one looks for both a difference in color and shape (i.e. green and **T**-shape are the conjunction). For another example, consider searching for the black rectangle in Figure 1.2A versus searching for the same rectangle in Figure 1.2B.



Figure 1.2 Visual search. In (A), the vertical dark rectangle pops-out because it is a feature singleton, the only one that differs from the distractors in terms of color. Pop-out is operationally defined as the relative independence of reaction time to set size (number of distractors) as given in (C). Notice the flat slope. In (B), we have a conjunction search where we must identify the vertical black rectangle. Here, visual search is harder and reaction time varies with set size as graphed in (D). Reprinted from S. P. Vecera and M. Rizzo (2003) "Spatial Attention: Normal Processes and their Breakdown." Neurologic Clinics 21: 575–607 with permission from Elsevier.

The relevant measure in these experiments is reaction time (RT), namely how long it takes the subject to report that the target is present or absent. Two basic findings are noteworthy. First, Treisman reported that in feature search, RT does not vary with set size (see Figure 1.2C). Here, the target seems to "pop out" from the display regardless of the number of distractors. Note that "pop-out" can describe the phenomenology (the object just seemed to pop out), but in the psychology of attention, it refers to the behavioral effect of constant reaction time despite increase in set size, as in Figure 1.2C. Second, in conjunction search, RT does vary with set size (see Figure 1.2D). These and other results led Treisman to propose a two-stage model. In feature search, processing of features happens concurrently or in parallel without capacity limitations. The target pops out because it is a singleton on one of the feature maps, namely a unique instance within that map (e.g., the color map might have a single red feature in a sea of green). In conjunction search, processing is non-parallel, a serial deployment of attention to one object at a time. To invoke a common metaphor, in conjunction search, focused attention operates like a moveable spotlight illuminating a subset of targets at a time.⁶

Treisman (1988) later modified FIT by postulating a master map of locations at an early stage in processing that serves as the target of focused (spatial) attention (see Figure 1.3).

Two points are worth highlighting. First, Treisman takes focused attention to have the function of binding features, and, in that way, construes attention as selection for object representation and thus for conscious awareness of objects.⁷ This explicit connection to conscious awareness provides a distinctive conception of attention (see Chapters 4–6). Second, visual search tasks suggest the possibility that there are two types of attention, one involved in pop-out, the other more like a scanning spotlight. This possibility can also be seen in the final experimental paradigm to be discussed in this chapter, namely, the Posner Spatial Cueing paradigm.

1.6 The Posner spatial cueing paradigm

One aspect of attention that Helmholtz (Helmholtz 1896) demonstrated is that attention can be deployed covertly in a way that is sensitive to spatial location. To shift attention, one does not need to move a relevant sensory organ. In the visual domain, one can overtly attend to something by moving one's eyes to it. Such overt attention is plausibly present in other modalities: I can optimally orient my ears to a sound by moving my head; I can reach for an object pressing on my back; I can move closer to sniff


Figure 1.3 Later version of Feature Integration Theory. Redrawn from Treisman (1998). The spotlight of attention focuses on a spatial location in a spatial map that does not code features. Features, rather, are coded in separate feature maps, which give information that the feature is present (the 'flag' in the feature map) and information about the location of the feature. Attention to a location then selects certain features to be bound in an object representation, which can then be compared with stored information or used in other tasks.

something; and I can swish wine in my mouth. Nevertheless, I need not move a part of the body to shift attention, something vividly brought out in the cocktail party effect or looking out of the corner of the eye: I can surreptitiously listen to the more interesting conversation behind our group, even as I feign interest in our conversation, or I can keep my eyes on you while visually attending to something else. Should one then understand there to be two kinds of attention, overt and covert? I think the simplest position is to understand that the movement of a sensory organ is sometimes generated to serve attention, and where it is, attention is overt. There aren't, then, two kinds of attention but rather a single capacity that can involve movement.

In the visual domain, the spatial cueing paradigm developed by Michael Posner has become a standard test for the deployment of spatial (covert) attention, namely the selection of spatial location. The subject's task is to report the presence or the features of a visual target. The experiment begins with the subject looking at a screen on which a fixation point is presented and on which the subject must maintain fixation. The fixation point



Figure 1.4 A depiction of the Posner spatial cueing paradigm. (A) shows a direct cue which occurs to the left of the fixation "+". In this case, the target, an "*", occurs at the cued location. The cue is a valid cue (an invalid cue would have occurred to the right of fixation, where the object appears on the left). (B) depicts indirect cueing with an arrow pointing to the left, and hence serving as a valid cue (an invalid cue would have pointed to the right with the object appearing on the left). In all cases, there is a temporal interval between cue and target. Reprinted from S. P. Vecera and M. Rizzo (2003) "Spatial Attention: Normal Processes and their Breakdown." Neurologic Clinics 21: 575–607 with permission from Elsevier.

remains, say, for one second, at which point a cue is presented for 100 milliseconds (ms). After this, there is a temporal lag between the cue and the presentation of the target, the *cue-target onset asynchrony* (CTOA). Once the target appears, the subject makes a report.

There are two types of cues: a direct cue that appears at the target location, and an indirect or symbolic cue, such as an arrow, that indicates a distinct location. Valid cues correctly indicate target location, invalid cues incorrectly indicate target location, while neutral cues provide no information about target location. Within an experiment, cues are typically weighted more towards valid than invalid cues (e.g., about 80% valid-20% invalid, with some small amount of neutral cues, in Posner 1980). The relevant variable of interest can be either reaction time (RT) or response accuracy. The presence of the neutral cue allows comparison of valid versus invalid cueing.

The Posner paradigm has largely been applied to the visual domain, but it has also been deployed in audition (Spence and Driver 1994). What is consistently found is that there are advantages in reaction time and accuracy with valid cues over neutral cues: reaction times are faster and accuracy is higher. Similarly, there is a disadvantage when invalid cues are presented relative to neutral cues: reaction times are slower and accuracy is lower. The idea is that with a valid cue, the subject preemptively moves the



Figure 1.5 Standard effects in respect of reaction time in the Posner spatial cueing paradigm. Invalid cues are associated with a cost, namely, increased reaction time relative to neutral cues, while valid cues are associated with a benefit, namely, decreased reaction time relative to neutral cues. Adapted from Figure 2.6, Wright and Ward (2008), p. 20.

attentional spotlight to the target location with advantage in reaction time and accuracy whereas with the invalid cue, attention is misdirected and must move again to the actual target location, with concomitant cost in reaction time and accuracy.

The temporal differences between pop-out and serial search in the previous section and direct and indirect cueing in the current discussion might suggest two types of attention, or at least two different ways of deploying attention. In the empirical literature, the putative division in attention is often expressed as that between top-down versus bottom-up attention, although there are a plethora of dichotomies that are also used (see next section). Conjunction search and indirect (symbolic) cueing are often referred to as bottom-up attention, while pop-out and direct cueing are often referred to as bottom-up attention (or in some other equivalent terms; this classification is not universally accepted as will be discussed in the next section).

How are top-down and bottom-up attention different? Consider some relevant effects from Posner's spatial cueing paradigm (Carrasco 2011 also gives a summary of relevant differences in Section 3.1). First, recall the cue-target onset asynchrony (CTOA), i.e., the time between onset of the spatial cue and appearance of the target. In Posner's paradigm, direct and indirect cues differ in their facilitation of reaction time (RT), in that direct cues yield a maximum facilitation on RT at a CTOA of 100 ms (i.e., target appears 100ms after the cue), while indirect cues yield the maximum facilitation at a CTOA of 300ms. Second, the benefits of cueing with direct cues

are transient and decay fairly rapidly, while those of indirect cues are sustained. These observations suggest that there are different mechanisms underlying the effects of different cues. These differences are depicted in Figure 1.6:



Figure 1.6 This graph shows the time course of the benefit in performance of direct (peripheral) versus indirect (central) cues. In this case, the direct cues were not predictive of target location (i.e., were equally likely to occur at the target location or not). Notice that the peak effect for direct, peripheral cues occurs earlier than that for indirect, central cues. For direct cues, there is also inhibition of return, as if attention is repelled for some time from the original cued location. Symbolic cues have a more sustained effect in terms of reaction-time benefit. Reprinted from S. P. Vecera and M. Rizzo (2003) "Spatial Attention: Normal Processes and their Breakdown." Neurologic Clinics 21: 575–607 with permission from Elsevier.

Memory load seems to have different effects on direct versus indirect cueing. When subjects are asked to do a task that requires keeping items in working memory (i.e., increased memory load), there are no significant effects on cueing facilitation with direct cues (e.g., in RT), while with indirect cues, the level of facilitation drops off with increased memory load. Perhaps the differences are not surprising. Symbols would seem to require, at a minimum, additional processing of the symbol in terms of its semantic significance. To respond to a symbolic cue like an arrow, one must understand its conventional meaning as an indicator.

This suggests the possibility of different mechanisms in direct and indirect cueing. The top-down and bottom-up distinction, defined at the psychological level, does seem to correspond to a division in underlying networks in the brain. The discussion of attentional networks gained much impetus with the publication of Michael Posner and Steven Petersen's, "The attention system of the human brain" (1990), a work cited over 3500 times in the intervening years and recently revisited by them (Petersen and Posner 2012). Posner and Petersen identified three networks associated with functions commonly attributed to attention: "(a) orienting to sensory

events; (b) detecting signals for focal (conscious) processing, and (c) maintaining a vigilant or alert state" (1990, p. 26). In their original discussion, they emphasized that attention forms its own system separate from motor and sensory systems, that attention involves a network of anatomical areas in the brain and that these areas carry out distinct functions (ibid.).

In important imaging work, focusing on Petersen and Posner's proposed orienting network, Maurizio Corbetta and Gordon Shulman (Corbetta and Shulman 2002) later proposed

that visual attention is controlled by two partially segregated neural systems. One system, which is centered on the dorsal posterior parietal and frontal cortex, is involved in the cognitive selection of sensory information and responses. The second system, which is largely lateralized to the right hemisphere and is centered on the temporoparietal and ventral frontal cortex, is recruited during the detection of behaviorally relevant sensory events, particularly when they are salient and unattended.

(p. 201–2)

The relevant network is diagrammed in Figure 1.7:



Figure 1.7 Rough localization of the regions of the two attentional networks: the dorsal frontoparietal network (open circles) and the ventral frontoparietal network (circles with gray shading). The former includes the intraparietal sulcus (IPS), the superior parietal lobule (SPL), and the frontal eye field (FEF); the latter includes the temporoparietal junction (TPJ) including the inferior parietal lobule (IPL) and superior temporal gyrus (STG), and the ventral frontal cortex (VFC), which includes the inferior frontal gyrus (IFG) and middle frontal gyrus (MFG). The IPS-FEF network plays a role in both top-down and bottom-up attentional processing; the TPJ-VFC network is involved in bottom-up attentional processing, including circuit-breaking in attentional capture. Frontal Lobe (FL); Occipital Lobe (OL); Temporal Lobe (TL); central sulcus (CS). This map of the attentional network is derived from Corbetta and Shulman (2002), figure 7. The figure is reprinted in adapted form from M. Behrmann, J. J. Geng, and S. Shomstein (2004) "Parietal cortex and attention." Current opinion in neurobiology 14: 212–217 with permission of Elsevier. Figure kindly provided by Marlene Behrmann.

The dorsal frontoparietal network is characterized as involved in the control of top-down attention. In Corbertta and Shulman's conceptualization, this top-down network generates and maintains an attentional set, namely "representations involved in the selection of task-relevant stimuli and responses" (202). It influences perceptual processing so as to serve current task demands, and in that way is sensitive to one's goals. On the other hand, the ventral frontoparietal network plays more of a role in bottom-up processing. Among its functions, this network serves as a circuit breaker. That is, certain salient stimuli, such as a loud sound, not only need to attract attention, but also stop other cognitive processes so that the subject can focus on the sudden stimulus. Note also that the two networks do not operate independently: while the dorsal network was recruited under all task conditions Shulman and Corbetta investigated, under bottom-up conditions, the ventral network was also recruited (Shulman and Corbetta 2012, 114). So, bottom-up and top-down attention seem to share some of the same neural substrates, but also differ in their neural substrates. The next chapter will return to the question of the neural implementation of attention, but the current task is to more critically scrutinize the conceptual contrasts that have been used to characterize attention.

1.7 Divisions of attention

This section considers some common ways of dividing attention:

- Top-down versus bottom-up
- Endogenous versus exogenous (cf. intrinsic versus extrinsic)
- Goal-directed versus stimulus-driven
- Controlled versus automatic
- Voluntary versus involuntary.

How should one understand these concepts so as to fruitfully invoke them in a theory of attention? Alan Allport has made the following observation:

In general, despite the ingenuity and subtlety of much of the experimental literature that has been devoted to these two enduring controversies [early versus late selection, and the idea of automaticity and control in processing, to be discussed in this section], the key concepts (selection, automaticity, attention, capacity, etc) have remained hopelessly ill-defined

30 THE PSYCHOLOGY OF ATTENTION

and/or subject to divergent interpretations. Little wonder that these controversies have remained unresolved.

(Allport 1993, 188)

For the concepts listed above, it is not hard to find papers where most of them are used, often in the same sentence. It is also not hard to find them being understood or applied in different ways between different papers. These notions are presumably technical terms but are never rigorously defined. No wonder Allport thinks there is muddle. Clarity requires definitions, and I shall provide definitions for what I think are the central notions: topdown versus bottom-up, and control versus automatic. Necessarily, the proposed definitions will involve some stipulation, but dissatisfied theorists are asked not to nay-say but to present concrete alternatives.

It is important to be clear that these terms apply to the subject. Thus, it is a psychological subject who exhibits top-down, endogenous, goal-directed, controlled, or voluntary attention. This leaves open other applications of these terms to the brain. For example, theorists speak of a brain region as exerting top-down influence on another region. This is a different use of "top-down" that can be perfectly appropriate, but it ascribes the relevant processing not to the psychological subject, but to a part of her. This recalls my earlier emphasis on the personal versus the subpersonal: some topdown effects are personal, as in attention; others are subpersonal, as in interactions between brain regions. It is no objection to the definitions to be given that they do not apply to interactions between brain regions. They are not intended to describe those interactions.

Let us begin with an initial proposal for top-down versus bottom-up, as much early work on attention divided mental processing into stages. Focusing on perceptual attention, if one thinks that perceptual processing forms the bottom of a processing hierarchy, then for S as subject and X as target:

S's attention is **top-down** if and only if *S*'s attention to *X* involves the influence of a non-perceptual psychological state/capacity for its occurrence.

S's attention to *X* is **bottom-up** if and only if *S*'s attention to *X* did not involve a non-perceptual psychological state/capacity for its occurrence.

An intuitive case of top-down attention is where a subject intends to pay attention in a certain way, say, to focus on a specific object. Thus, the subject's attending to that target occurs because the subject intends to attend to it. The selection at issue occurs because of the deployment of intention, a non-perceptual psychological capacity. Where perceptual attention happens without needing the influence of non-perceptual psychological capacities, attention is then bottom-up. This covers the intuitive cases when attention is captured by what one perceives, such as a loud bang. What this influence ultimately comes to, mechanistically speaking, is a matter for empirical research. Note that the definitions assume that one can divide the mind into systems, and in particular, between perceptual and non-perceptual systems. It is a good question whether one can adequately do this, an issue that must be set aside. In addition, any non-perceptual psychological system counts as part of the "top". So, motor influences on perceptual selection would count as top-down. Again, this is stipulative, but it allows for clarity.⁸

What of exogenous versus endogenous sources of attention (sometimes also intrinsic versus extrinsic)? It is not clear how this distinction differs from the previous. For example, Marisa Carrasco (2011) writes:

The [endogenous system] is a voluntary system that corresponds to our ability to willfully monitor information at a given location; the [exogenous system] is an involuntary system that corresponds to an automatic orienting response to a location where sudden stimulation has occurred. (p. 1488)

Carrasco further points out that endogenous attention is sometimes spoken of as sustained attention while exogenous attention is spoken of as transient attention (recall the temporal properties of direct and indirect cueing discussed in the previous section and depicted in Figure 1.6). It is not clear, however, that the exogenous/endogenous dichotomy comes to anything more than the top-down, bottom-up contrast. For current purposes, I shall treat them as equivalent.

Bottom-up attention maps onto stimulus-driven attention, if one thinks of the stimulus as always first dealt with by perceptual systems. Stimulusdriven attention is often contrasted with goal-directed attention, but on any plausible account, goal-directed attention is only one type of top-down attention. Goals are, presumably, embodied in intentions or plans, but the account of top-down attention allows for all sorts of non-perceptual influences: memory, expectation, emotion, values, and habits.⁹ Accordingly, the contrast between stimulus-driven and goal-directed attention is not exhaustive. There are non-stimulus-driven forms of attention that are also not goaldirected, say my preference for chocolate over fruity candies that leads to my attending to chocolates in a candy store even if I am not intending to buy any candy. The stimulus-driven versus goal-directed contrast falls short of taxonomizing attention.

Things get murky with control versus automatic attention, on the one hand, and voluntary versus involuntary attention, on the other. The reason is that these notions point to agency. After all, one speaks of a person as being in control or doing something automatically, or of her doing something voluntarily or involuntarily. So, understanding these contrasts requires understanding action, a notion even more challenging than that of attention. I propose to focus on control versus automatic. The voluntary versus involuntary distinction is difficult for it either suggests a kind of agency, such as free agency or agency that involves the will in a specific way, or connotes a characteristic sort of consciousness, something that might be tied to felt effort or a sense of activity. Since the voluntary is tied up with further complex phenomena, it is not likely to help draw clear boundaries in attention.

One can, however, explicate automaticity and control more clearly using the notion of an intention, a goal-representational state. In psychology, the ideas of Richard Shiffrin and Walter Schneider (1977) greatly influenced subsequent discussions of the control-automaticity dichotomy. On automatic processes, they wrote:

an automatic process can be defined ... as the activation of a sequence of nodes with the following properties: (a) The sequence of nodes (nearly) always becomes active in response to a particular input configuration, where the inputs may be externally or internally generated and include the general situational context. (b) The sequence is activated automatically without the necessity of active control or attention by the subject.

(2)

This proposal connects automaticity to the absence of control (or attention) by the subject. What then is control on their conception? "A controlled process is a temporary sequence of nodes activated under control of, and through attention by, the subject" (ibid.). Ignoring the circularity in their definitions, one can take Shiffrin and Schneider as defining automaticity in terms of the absence of control, while control is tied to attention. In contrast, I propose to explicate the notion of control in terms of the role of intention. Here's the basic idea in a slogan: control in attention is attending as you intend. Representations of a subject's goals are embodied in the subject's intentions, namely representations of a plan of action. These plans and their corresponding mental states can be expressed by reports such as I intend to do X or I will do X. Following Elizabeth Anscombe (1957), philosophers have noted that while actions can be described in many ways, only certain descriptions capture how agents conceive of their actions. That is, they are revealed as intentional only under certain descriptions. Thus, while Gavrilo Princip intended to assassinate Archduke Ferdinand, he did not intend to precipitate the First World War, even if the assassination was identical to the precipitation of war. Those descriptions describe the same action (Davidson 1980).

Control in attention is attention as one intends. Control also implies the absence of automaticity, or automaticity is the absence of control, as Shiffrin and Schneider emphasized. At the same time, if one looks at processes that are controlled, say deliberate actions, one also finds automaticity. You might intentionally throw a ball, but many aspects of your throwing such as its kinematics, the way your joints rotate, and the sequence of movements in your arm are automatic. You don't intend to throw with that speed, rotation or sequence of movements, but your intentional throwing wouldn't be what it is without them. So, despite the contrast between control and automaticity, intentional activities often involve both. How can this be?

Elsewhere, I have argued that one can define automaticity as the absence of control and allow for actions to be simultaneously controlled and automatic only if one relativizes automaticity and control to properties of the process. That is, one speaks of control of a process in respect of a specific feature of that process, and likewise for automaticity. Accordingly, automaticity entails the absence of control, yet a process can be both automatic and controlled in respect of different properties. I shall not give here a detailed version of the analyses of control and automaticity (see Wu 2013a), but the following biconditionals capture the essential idea and will suffice for current purposes.

(C) S's attention to X is **controlled** relative to its feature F iff S's attention having F results from S's intending it to have F.¹⁰

Following Shiffrin and Schneider in defining automatic negatively as the absence of control, one derives:

(A) S's attention to X is **automatic** relative to its feature F iff S's attention having F is not due to control as per (C).

To see how this works, consider visual conjunction search tasks where the target is a red letter **E**. Where the subject attentionally selects a red **E**, her attention's having the feature of selecting a red **E** is controlled because it is precisely what the subject intends to do. Similarly, if the subject has her attention captured by a suddenly appearing stimulus, then attending to that stimulus is automatic because the subject did not intend to attend to that stimulus. In both cases, the relevant feature *F* is the subject's attention having the target that it has.

Here is an intuitive gloss of each definition. With top-down versus bottom-up, the key concern is how attention gets initiated, i.e., whether the subject is passive or active in that initiation. With top-down attention, the initiation of attention involves and is attributed to the subject due to some non-perceptual mental state including the subject's intentions. In bottom-up attention, by contrast, one can think of the stimulus as initiating attention, even if it disrupts a subject's current activities. On the other hand, think of the control versus automaticity distinction as concerned primarily with the shape of attention once it begins and with how the features of that process unfold: where attention is directed and in what sequence, how long it is sustained, to what specific features in the scene, and so on. Finally, it is worth pointing out that top-down and control are sometimes treated as equivalent; in our account, they are not.

Let us now relate the two central dichotomies.¹¹ Given the previous definitions, there are four categories:

- (1) Top-down, controlled attention;
- (2) Bottom-up, automatic attention;
- (3) Top-down, automatic attention;
- (4) Bottom-up, controlled attention.

The first two may not be that surprising, perhaps because one assumes that top-down implies control and bottom-up implies automaticity. In fact, this does not follow, for recall that the top-down/bottom-up distinction is tied to the occurrence of attention while the controlled/automatic distinction is tied to its features.

(1) and (2) are familiar categories. You tell me to follow the man in the fedora, and I attend to him. My attention to him is top-down and controlled. It is top-down because I initiate attention given my intention to follow your instructions, and I would not have done so otherwise. Further, it is controlled because my attention has the feature of being directed at

that person as a result of my intention to keep my eyes on that person. In general, intentional forms of attention fit with (1). In a case of (2), a loud continuous sound pulls my attention to it. It thus looks like this capture of my attention occurs independently of any top-down influence, so it looks to be bottom-up. Moreover, given that I don't have any relevant intentions, many of the features of attention might be automatic although perhaps not for long. I hear the sound and subsequently intend to figure out where it is coming from, so attention thereby takes on a controlled aspect. It is sustained according to my intentions. The phenomenon of pop-out in visual search might also seem like a case of bottom-up, automatic attention, but this is controversial.

What of top-down, automatic attention? This seems an odd category, but consider the following experiment by Alfred Yarbus (1967). Yarbus presented his subjects with a painting of a homecoming scene and asked them to perform three tasks: (i) remember what the people in the picture are wearing; (ii) remember the location of people and objects; and (iii) estimate how long the visitor has been away. He then tracked their eye-movements (overt attention) while they carried out his instructions and noted the following patterns:



Figure 1.8 Yarbus asked subjects to perform a variety of tasks in relation to I. P. Repin's "Unexpected Visitor" (A). He monitored their eye movements as they visually interrogated the painting in order to perform his tasks. For example, panel (B) indicates the eye movements in response to the command to remember the clothes worn by the people; (C) to remember the position of people and objects in the room; and (D) to estimate how long the visitor has been away from the family. Material from Yarbus (1967), p. 174, figure 109 with kind permission from Springer Science+Business Media B.V. This figure reproduced from "Eye Movements and the Control of Action in Everyday Life" M. F. Land (2006) Progress in Retinal and Eye Research 25: 296–324 with permission from Elsevier.

What is striking is that the patterns of eye movements make sense given the subjects' more abstract intentions to carry out Yarbus's instructions. For example, asked to remember the clothes, the subject intentionally looks at each figure. This intention need not be an intention to move one's eyes in any specific way, but the resulting pattern of eye movements in panel B is intelligible given the intention in question: the eyes gravitate around the people without spending time on the objects. Attention in the form of eye movements, overt attention, tracks the intention even if the intention is not to move the eyes in a specific pattern. The specific pattern of eye movements happens automatically and is not itself intended. Moreover, the pattern is a feature of overt attention, one that is not represented in the content of the intention. As the pattern is not controlled, it is automatic. At the same time, overt attention with this pattern would not have occurred without the subject having the requisite goal, so attention is top-down. Notice that when one toggles the subject's intention by presenting different tasks, the pattern of eye movement changes.¹² So, intentions are involved in the occurrence of overt attention with a characteristic pattern. While the idea of top-down automatic attention might seem contradictory, it is not. That we can categorize the phenomenon Yarbus observed suggests that the initial analysis is theoretically useful. This is a sign that the definitions are on the right track.

It seems likely that no process instantiates (4) since the causal processes imputed by each dichotomy operate at cross purposes: bottom-up attention requires a stimulus-driven initiation independent of any intentions, but control requires the influence of an intention. At best, attention might be bottom-up and automatic but quickly *becomes* controlled once intentions kick in to sustain attention to the stimulus. In any event, I want to conclude the discussion of the conceptual issues by returning to category (2): bottom-up, automatic attention when attention functions like a circuit breaker. This seems like an obvious, familiar category. Yet like (4), there are questions whether (2) is ever instantiated.

The previous section noted the difference between direct and indirect cueing. Richard Wright and Lawrence Ward (2008) suggest the following:

Researchers can choose to study either voluntary or involuntary orienting, depending on whether they use symbolic or direct location cues ... Symbolic location cues initiate attention shifts in a fundamentally different way than direct location cues. The former are meaningfully associated with a particular location and therefore must be interpreted by an observer in order to be used. For this reason, the initiation of an attention shift

by a symbolic cue is *goal-driven*. The observer processes the location information conveyed by the symbol and, on this basis, develops a computational goal for carrying out the task ... Direct cues, on the other hand, produce their effect by virtue of being physically close to the target location ... No cognitive interpretation of direct-cue meaning is required and, instead, attention is captured by the onset of the cue. For this reason, the initiation of an attention shift by a direct cue is *stimulus-driven*.

(21-22)

It is natural to take the direct cue as capturing attention and in that way independent of goals. But is it goal-independent? Bradley Gibson and Erin Kelsey (1998) suggest that the influence of the direct cue is goal-directed (p. 699): "stimulus-driven attentional capture may be caused by goal-directed processes." How can this be?

In discussing Feature Integration Theory (FIT), I noted that feature singletons (i.e., features that are unique within a feature map such as a red shape in a sea of green shapes) seem to pop out. It would be natural to characterize pop-out as the capture of attention as occurs with auditory attention and loud noises.¹³ John Jonides and Steven Yantis (1984; 1988) have argued, however, that most cases of pop-out in the attention literature are not genuinely bottom-up, automatic capture of attention but depend on the subject's goals. Hence, they are top-down! Consider the visual search tasks discussed above when the target seems to pop out. To undertake the task, you have to follow task instructions, say to locate a green **T**. Yet in intending to locate a green **T**, you've set yourself to complete a specific task. Locating that target is your explicit goal. The target you intend to locate is precisely what pops out. Again, it is top-down, and your locating it reflects attentional control.¹⁴ Of course, there are automatic elements. What you don't control, and hence what is automatic, is when you locate the target. That you locate it in a way independent of set size reflects the automaticity with respect to when you locate it (reaction time is the same). Nevertheless, attending to the T is top-down and controlled. My definitions show how one can consistently and clearly apply the concepts of top-down, control, and automaticity to the same phenomenon.

It is striking that the original pop-out effects might in fact be top-down and controlled rather than bottom-up and automatic. Indeed, Charles Folk, Roger Remington, and colleagues (1992) claim that there are no pure cases of bottom-up, stimulus-driven attention (for a methodical review of the issues and experimental evidence, see Burnham 2007). One can pose the issue as a challenge: is attention ever independent of the goals of the perceiver?¹⁵ The claim is that goals have a pervasive influence on attention. Still, it is hard to accept the claim that there is *never* attentional capture contrary to one's goals. Consider being engrossed in a performance of Beethoven's Ninth Symphony. Just before the climactic moment of the famous chorus, I pinch your shoulder. This is quite annoying, of course, since it breaks your concentration on the music, but it also does seem to be a compelling case of tactile attentional capture. Attention is devoted to auditory experience, as you listen to the music. Your intention for the past hour has been to listen, your attention has been focused on the music. There do not seem, then, to be any goals where tactile inputs are relevant. This is, of course, an anecdote, but a prima facie compelling one.

The dichotomies discussed in this section have been deployed for a long time in the study of attention, and they are well entrenched in psychological vocabulary. At the same time, there is something casual and slippery about their use that needs to be avoided once they are deployed in serious theory building. The proposals I have given provide concrete accounts of what these dichotomies come to. I suggest that barring any other concrete definitions (and there are none in the literature that I am aware of), theorists should start with the ones presented here.

1.8 A sufficient condition for attention: selection for task

There is a central idea towards which all the theories, paradigms, and conceptual dichotomies discussed thus far gravitate: the notion of a task. For example, the Load Theory of Attention argues for the task-dependence of where attention acts in perceptual processing. Further, the subject's goals pervasively influence attention, so much so that some theorists have questioned whether there is attention without the influence of goals. Finally, three specific experimental paradigms have been central to the psychological study of attention: dichotic listening, visual search, and spatial cueing. In each of these, a well-defined task structures the experiments. Given the centrality of tasks, might appeal to it provide a way to answer the growing skepticism to explaining what attention is?

A well-defined experimental task establishes conditions such that when they are fulfilled, the experimenter is confident that the subject has deployed the capacity the experimenter is studying. Specifically, where the subject has followed task instructions and correctly performed the task, the experimenter can be

confident that the capacity in question has been deployed. Consider then studies of attention using verbal shadowing in dichotic listening paradigms. Where the subject correctly shadows the verbal stream assigned, the experimenter can be confident that the subject is attending to that stream, using the sounds in that stream to inform verbal response. Next, consider the use of reaction time to measure task performance in visual search and spatial cueing. The reaction at issue in both experiments is target detection, and reaction time reflects the temporal properties of attention in serving that task. Given that subjects perform that task, namely, producing a judgment about the target's presence or absence, this performance is a sign that the subjects have been attending to the relevant target, using it to render a judgment. This can also be discerned by looking at eye movements during the task. Obviously, when subjects are not doing the task, say when they twiddle their thumbs or continuously get things wrong, this is evidence that they are not appropriately selecting the relevant target and are being inattentive.

In the three experimental paradigms that I have discussed, it is clear that, for each, there is a well-defined target, reaction to which requires selection of that target to inform the response, whether tracking a conversation in verbal shadowing or examining targets in target detection. As these experiments are used to probe attention, there is a general assumption that all experimenters on attention hold in using these paradigms:

Empirical Sufficient Condition for Attention (S_{emp}): Subject S perceptually attends to X if S perceptually selects X to guide performance of some experimental task T, i.e., selects X for that task.

Where the subject selects some target to guide their response in carrying out an instructed task, then the subject's selecting of that target is sufficient for the subject's attending to that target. Notice that the condition introduces a variable for the targets of attention and selection, targets that can be information, locations, features, or objects. Thus, dichotic listening and visual search involves the tracking of features and objects, say visible and audible entities and their properties, while Posner's spatial cueing paradigm tests, in part, for attention to locations. In what follows, the focus will be on locations, features (properties) and objects as targets of attention.

One might wonder why not just say that the subject's selecting X is just the subject's attending to X. This would, however, require that selecting X is a necessary condition as well, but that is controversial and more difficult to establish (I shall try to establish it in Chapter 3 by defending a selection for action account of attention). For current purposes, the sufficient condition provides an answer to the skepticism noted in the Introduction. For all their doubts concerning answering the metaphysical question, theorists of attention have done much interesting and important experimental work on attention. Furthermore, an assumption in their experimental practice, namely the empirical sufficient condition, provides a shared condition on attention that is relevant to the metaphysical question. Of course, not all sufficient conditions for a phenomenon are informative as to its nature. That someone wins the Electoral College in the U.S. presidential election is sufficient for their becoming U.S. president, but winning the Electoral College doesn't illuminate what a president is. The interest of the empirical sufficient condition is that it begins to flesh out talk of attention as selection by drawing on an assumption in experimental work on attention.

On reflection, this should not be surprising. Any experimentalist who wants a subject to direct attention knows how to do it, namely by having the subject perform specific tasks with respect to a target. That is, if the experimentalist wants to ensure that the subject attends to some X, then the experimenter designs a task where X is task-relevant and where X must be used to perform the task. To study attention, one needs to know how to manipulate it and to keep track of it. A well-designed experimental task is precisely one that creates conditions such that one can do so, and this is just manipulating the subject's task performance by manipulating what the subject must selectively respond to. The empirical sufficient condition then identifies a widely held assumption in empirical work on attention that can serve as an initial foothold in the face of skepticism about what attention is.

1.9 Summary

This chapter began by highlighting five basic questions, and in discussing the fruits of psychological research in the last 70 years, uncovered many answers, especially to the properties question. Of note, the properties of attention seem to point to two forms of attention. These forms have different temporal profiles regarding when they exert their greatest effect and how long they last, they have different dependencies on memory, and they seem to call on overlapping, but different, neural networks. As a result, it has been natural to divide attention, and this has led to two salient divisions when characterizing attention: top-down versus bottom-up and control versus automaticity. I have provided an analysis of the resulting dichotomies of attention, and highlighted an interesting category of attention, namely a top-down, automatic form illustrated in Yarbus's eye-tracking experiments. This work provides a more detailed picture of our capacity to attend.

The function question has also received some interesting answers that in turn suggest a possible answer to the metaphysical question. I have canvassed conceptions of attention as a filter and a spotlight. On the one hand, Broadbent emphasized filtering to explain the role of attention in perceptual processing, namely in selecting relevant information for further work-up. On the other hand, spotlighting suggests a phenomenal aspect to attention and has an echo in Treisman's talk of attention as selecting features to bind for conscious awareness of objects (she spoke of attention as "glue" for feature binding). I will examine the phenomenal conception of attention and attention's relation to consciousness in later chapters (Chapters 4-6), but the current discussion revealed an interesting property of attentional filtering, namely that the stage at which attention acts, namely, early versus late, seems to be task-dependent as hypothesized by the Load Theory. Rather than attention being tied to a specific stage in processing, perhaps it is tied to the task that the agent performs. Indeed, in the last section, I argued that an implicit assumption in experimental paradigms used to probe attention is that a subject's selecting an item to inform task performance is sufficient for the subject's attending to that item. This then points to another possible answer to the metaphysical question: might attention be selection for task, indeed. for action?

Suggested reading

Mole (2011) and Hatfield (1998) discuss psychological work on attention pre-1950, and Tsotsos (2011, chap. 1) presents a nice overview as well. A succinct account of the psychology of attention from the 1950s onwards is provided for in Driver (2001). Pashler (1998) is a monograph discussion of similar terrain. Relevant recent review articles on the psychology of attention can be found in Posner (2011). Lavie and Tsal (1994) provide a discussion of the Load Theory of attention in light of the early versus late selection debate. On visual search and Feature Integration Theory, Treisman (1988) provides an overview, while Wolfe (1994) provides an update of his version of visual search. Wright and Ward (2008) provides an well-known critique of 25 years of attention research, while Carrasco (2011) provides a

more recent overview of the last 25 years of attention work. On information theory, Weaver (1949) provides an accessible overview (but see also Appendix).

Notes

- ¹ Subjects themselves can exhibit selectivity without attention. For example, you might value fine wines, while I value fancy cars. We thus have selective tastes or values, but none of this amounts to attention.
- 2 Where exactly one puts the bottleneck depends in part on a specific model of the different stages of perceptual processing. Specific debates about early or late selection in perceptual processing will ultimately be relative to concrete models of perception. I will, however, operate with a simple conception of perceptual processing as consisting of a basic feature-processing stage and then a later, higher-level categorical-processing stage.
- 3 One could argue that the switch occurs because shadowing involves tracking a specific object, say the voice, rather than a specific channel, say the right ear. Or perhaps subjects confuse the two objects of attention, beginning first with spatial attention to the right ear but then focusing on the voice. Setting this aside, Treisman's experiment was one among many experiments that pushed psychologists to late-selection theories.
- 4 In contrast, Navon and Gopher (1979) proposed that processing uses only what it needs up to the limits of what is available.
- 5 What counts as a *basic* feature? For discussion, see Wolfe and Horowitz (2004).
- 6 Aside from filters and spotlights, another common metaphor for attention is a zoom lens (Eriksen and St James 1986). The claim that focused attention is serially deployed can be challenged (Wolfe 2003). In particular, one can develop coherent parallel processing models that mimic the RT effects in conjunction search. Moreover, there are experiments suggesting that feature search can involve attentional scanning (Treisman and Gormican 1988) and that conjunctive targets can pop out (Nakayama and Silverman 1986).
- 7 This gestures at the so-called "binding problem" in cognitive science. For an exchange that brings out some of the central issues, see (Di Lollo 2012) and (Wolfe 2012). For an overview, see (Treisman 1996).
- 8 An alternative is to define top-down as attention to *X* resulting from the processing influence by some *cognitive* system. Still, I think conative states like desire can direct attention, so I opt for the definition in the text.

- 9 Hutchinson and Turk-Browne (2012) note that memory can influence attention, but that this influence is not captured by the stimulus-driven versus goal-driven dichotomy. The definition of top-down attention given earlier makes room for memory-guided attention.
- 10 The result has to be caused in the "the right way," a caveat that points to the problem of deviant causal chains in the philosophy of action. I assume that those problems can be solved.
- 11 The bottom-up and stimulus-driven characterizations are overly focused on the perceptual cases, but there are cases where involuntary thoughts and memories can capture attention, "automatically" as one might say. To allow for such cases, the definitions would need to be modified so that different mental states can serve as the input into attention. I leave that as an exercise.
- 12 It is worth noting that the Yarbus result has recently been revisited by Greene et al. (2012) who were unable to replicate the intelligibility of eye movement result using different images. This matter should be experimentally revisited (e.g. one should try to replicate the Yarbus result with the original image). One can make the claim in the text a conditional: if the Yarbus result is correct, then there seems to be a form of top-down automatic attention.
- 13 In fact, pop-out should be specifically understood as what happens with respect to a feature in visual search where the time to locate that feature is not sensitive to distractor set size. It need not imply attentional capture.
- 14 Jonides and Yantis (1988) found that when the subject is told to actively ignore color (or other properties), then colored targets do not pop out, and a plot of RT to set size now shows a non-zero slope, i.e., RT now increases with set size. Jonides and Yantis make the further claim that abrupt onset singletons, e.g., a flashing stimulus like the direct cues in the spatial cueing paradigm, almost always capture attention (though in one of their experiments, focused attention on a specific location abolishes capture by abrupt onset stimuli; see Folk, Remington, and Johnston 1992 for discussion of this in an exchange with Yantis 1993). On missing things when attending elsewhere, see the discussion of inattentional blindness, Chapter 5.
- 15 Folk, Remington and Johnston (1992) propose the Contingent Involuntary Orienting Hypothesis:

under conditions of spatial uncertainty, involuntary shifts of attention to a given stimulus event ... will be contingent on whether that event shares a feature property that is critical to

44 THE PSYCHOLOGY OF ATTENTION

the performance of the task at hand. Specifically, involuntary orienting of attention will occur if the event shares the critical property and will not occur if it does not.

(1992, p. 1032)

Jan Theeuwes (2010) has argued that the initial selectivity seen during visual processing is always bottom-up, and he argues that these effects happen in a temporal window that is too narrow for top-down modulation to occur. This issue continues to be debated (see Theeuwes 2010 as target article, and responses that follow). I am grateful to Barry Smith for prompting the example that follows.

2

THE NEUROSCIENCE OF ATTENTION

2.1 Introduction

Psychological research on attention has yielded a wealth of data and models. In responding to that work, it is easy to reach for metaphors of attention as a filter, as a spotlight, or as something else. Metaphors have the virtue of being compelling yet the vice of being uninformative. Can matters be made more concrete? One reaction from neuroscience is to note that, while the behavioral data generated in psychology suggest many ways that attention might work, in the end, attention is implemented in a brain. It is the brain and specifically neural mechanisms that ultimately realize the capacities isolated in the various attention paradigms discussed in Chapter 1. Understanding the brain will be fundamental to understanding attention. If there is a question as to what attention is, then the answer will be found in the workings of the brain. That is, one should pursue a *neural mechanism-based metaphysics* of attention. Fundamentally, attention is a brain process, or so a neuroscientist might aver.

The nature of mechanisms, either as an ontological category or as a type of explanation, is a complex matter that must be sidestepped in favor of a more general understanding.¹ I will construe mechanisms as certain kinds of answers to a how-question: how does X happen, work, or how is X implemented? Thus, how does attention happen, work, or how is it implemented? One type of answer is a specification of a process that

explains how X works, happens, or is implemented. Accordingly, Feature Integration Theory presents a mechanism of object representation where attention plays a crucial role and filtering describes a step in perceptual processing where attention acts in a certain way. So, some mechanisms are abstractly specified. In this chapter, I move from behavioral data and the more abstract models they suggest to the workings of neurons and brain regions to implement attention. In part this is to ground those abstract models in concrete neural mechanisms and activities. The neuroscience of attention will reveal echoes of filtering and spotlighting in the activity of neurons, tantalizingly suggesting ways that psychology and neuroscience might be united. It provides a way to unpack and examine a not uncommon reductionist approach in the history of cognitive science that identifies attention with a specific brain process. Might attention be a brain process? This chapter examines the case for a positive view.

Section 2.2 begins with an explanatory framework due to David Marr as a way of organizing data from neuroscience with an eye to connecting neural mechanisms to the more abstract models from psychology. I then highlight two basic features of visual neurons in Section 2.3: their generation of electrical spikes and their possession of a receptive field that determines what stimuli they respond to. Section 2.4 describes two changes to the activity of a neuron under conditions of attention, changes that seem to point to spotlighting and filtering: the amplification of a signal (gain modulation) and the remapping of the receptive field. Section 2.5 zooms out a bit to consider two general mechanisms that explain the phenomena discussed in Section 2.4: biased competition and divisive normalization. A third general mechanism for spatial attention that ties it to action is considered in Section 2.6, namely the influential Premotor Theory of attention. Evidence against that theory will be considered. An assessment of a neural mechanism-based metaphysics of attention is given in Section 2.7 and returns a cautiously skeptical answer. Indeed, the attempt to identify neural mechanisms for attention seems to rely heavily on the empirical sufficient condition. This suggests that the connection to task might be more fundamental in the order of explanation of attention.

2.2 The Marrian framework

It will be useful to provide an explanatory scheme to frame work in the neuroscience of attention. The scheme is David Marr's who begins his book, *Vision*, as follows: What does it mean, to see? The plain man's answer (and Aristotle's, too) would be, to know what is where by looking. In other words, vision is the *process* of discovering from images what is present in the world, and where it is.

(Marr 1982, 3)

Marr's opening question echoes two of the basic questions regarding attention. On the one hand, Marr's question suggests the metaphysical question regarding vision: What is vision? On the other hand, it suggests the function question: What is the purpose of vision? The answer from the plain man provides an answer to both. To see, Marr claims, is to know what is where by looking, and the point of seeing is to yield knowledge of the external world. Accordingly, to specify this function of vision is to begin to explain what vision is.

Marr assumes that the plain man's answer is to be explicated by his own: vision is an information process. In the sentence that follows the quoted passage, Marr continues: "Vision is therefore, first and foremost, an information-processing task ... " (2, my italics). By "task", Marr has in mind a process ultimately to be investigated by neuroscience, although the functional question has an explanatory priority. Marr thus moves from what philosophers call functional role, in this case the purpose or goal the capacity serves, to the realizer of that role, namely, the implementation of the function. For example, consider the role of being a can opener and the different physical devices that can realize can openers. Similarly, one can focus on the neural processes that implement knowing by looking. This links the neuroscience of vision to the psychology of vision.

Let us focus on the three-tiered explanatory framework that Marr proposed to explain vision, with the aim of adapting it to organize empirical work on attention. Marr's basic idea is that an adequate explanation requires approaching the phenomenon from multiple levels. These are "the different levels at which an information-processing device must be understood before one can be said to have understood it completely" (24; my italics). Indeed, Marr's own work in neuroscience led him to prioritize the highest level in his explanatory hierarchy, what he called the *computational theory*. Here, the theorist specifies what the phenomenon is for, what role it plays. The computational theory thus identifies the functional role of the capacity of interest. For vision, the plain man's answer gives the computational theory, and Marr reconceptualized this idea as the visual system's taking in information from a two-dimensional image corresponding to the stimulation of the retina (the input from looking) and generating a three-dimensional visual representation of the world (the output as putative knowledge).

The next two levels are the algorithmic/representation level and the implementation level.² In the former, the theorist specifies algorithms understood as procedures which, when carried out, compute the functions described in the computational theory and identify the representations on which the algorithms operate. For example, recall Treisman's Feature Integration Theory (FIT) where elements of feature maps are bound to yield object representations. One can think of FIT as a proposal for part of the algorithmic realization of vision as Marr conceives of it. Visual representations of the environment require visual representations of an object. In part, this involves binding by focal attention. So attention implements one goal in visual computation: taking representations of features and binding them to yield a representation of an object. Once algorithms and representations are specified, the goal then is to identify their physical (neural) implementation.

Note that the algorithmic level is the bridge between the functional specification of what a capacity is for at the computational level and the physical specification of the realization of that capacity at the implementation level. Furthermore, if mechanisms are what explains how things work, one can think of both the descriptions of algorithms, as well as of the neural implementation, as descriptions of mechanisms. That is, one can treat the lower two levels as providing a type of mechanistic understanding of the capacity described at the computational level, a way of explaining how that capacity works. The focus in this chapter will be on neural mechanisms, though some of the proposals to be discussed (biased competition and divisive normalization) are pitched more at the algorithmic level even though their proponents are thinking of neural interactions. There is also a wealth of computational work on attention that we will not have space to consider (for a nice overview and taxonomy of this type of work, see Tsotsos and Rothenstein 2011).

Marr's proposal suggests a way to organize attempts to explain what attention is. Indeed, I argued that all experimentalists investigating attention should endorse the following sufficient condition for attention: where S perceptually selects X to inform performance of task T, then S perceptually attends to X (in T). Focusing on the visual system, the relevant conditional is: if S visually selects X to inform T, then S visually attends to X. One can take as the relevant tasks activities such as visual search, filtering, spatial localization, and some of the tasks to be discussed below. But notice that

this gives an initial computational theory of attention. One function of attention is to serve performance of tasks of these kinds.³

Certainly, the current computational theory is incomplete since it provides only a sufficient but not necessary condition for attention. Still, it provides an empirically motivated starting point for relating empirical work in psychology and neuroscience. The idea is to begin with some defined task T that requires visual attention. This gives the computational theory for visual attention vis-à-vis T: visual attention is visual selection of some X that serves performance of T. One then canvasses proposals for the algorithms that implement selection of X for task T and the representations that are needed. For example, one can ask about the algorithms needed for selection of spatial locations in the Posner spatial cueing paradigm and the requisite representations of spatial location. Moving to the level of neural implementation, one can then ask how the brain implements such algorithms and representations. Of course, work at each level will mutually inform and constrain each other. Ideally, the product will be a complete explanation of one form of attention, namely selection of X for spatial cueing. This can be iterated for other experimental tasks, leading to a complete story across Marr's three levels for each task. One hope is that this investigation will identify commonalities in mechanisms, both computational and biological, across tasks. Completing this project, however, will not be easy, but the current point is to provide an organizing framework for any such investigation:

> **Computational Theory** Perceptual Attention = Selection of X for task T

> Representation/Algorithm Level Computational Mechanisms (filtering, binding) Implementing Selection of X for T

Implementation Level Neural Mechanisms Implementing Computational Mechanisms

Figure 2.1 A Marrian framework for attention consisting of three levels of explanation: the computational theory, the algorithmic/representational level, and the hardware (neural) implementation level.

What is gained from this perspective is a way to use some of the more abstract results from Chapter 1 to filter and organize the more concrete results to be discussed in the remainder of this chapter.

2.3 Two basic features of neurons

This section focuses on two fundamental features of neurons: their firing (spike) rate, which is a function of their electrical activity, and their receptive fields, a determinant of their electrical response. Mental function is driven by the electrical activity of neurons. Due to an electrochemical gradient that is actively maintained across its cell membrane, a neuron is able to generate an action potential, a rapid change in the electrical potential across the membrane. An action potential can be recorded as a spike or electrical discharge, and the generation of spikes, referred to as spiking or firing, is the way neurons carry information. Spikes can be detected and recorded directly using an electrode strategically placed on a neuron or multiunit electrodes that record spikes across a population of neurons. Neurons can be recorded in vitro, as when a sample of brain tissue is manipulated, or in vivo in recording from live animals, either awake or anaesthetized. Neural activity can also be monitored indirectly by different forms of neuroimaging such as functional magnetic resonance imaging (fMRI), electroencephalography (EEG), and magnetoencephalography (MEG).⁴ These provide a window on neural activity via changes in blood flow, electrical fields, and magnetic fields, respectively, changes that are correlated with neural activity. These methods allow for monitoring large brain regions, and their spatial and temporal resolution are accordingly much coarser than that of direct electrophysiological recordings. Such techniques can give a picture of activity across multiple brain regions.

An important measure of neural function is the firing rate, namely the number of spikes per unit time. Many neurons will spike spontaneously, giving a baseline firing rate for that neuron. Think of this baseline as the level of activity of the neuron when it is not being stimulated. Neurons receive inputs from other neurons or, in the case of neurons at the sensory organs, from external stimuli. Stimulation by these inputs can alter baseline firing rates, either increasing or suppressing firing rate. As experiments demonstrate, attention can also affect firing rates. The following are two types of plots showing changes in firing rate (the plots also show the difference of firing rate with and without attention, which is discussed in more detail in what follows):



Figure 2.2 Example of plots showing the spikes generated by a single neuron when a stimulus of a certain contrast was presented in its receptive field. In panels (A) and (B), which give nater plots of a train of spikes, each horizontal row of dots represents a single experimental trial, and each dot represents a spike. Time "0" indicates the onset of the stimulus that exhibits a contrast at 10%. The thick black line on the x-axis shows the duration of the stimulus. In (A), the animal directed attention outside of the neuron's receptive field. In (B), the animal directed attention within the neuron's receptive field. Note the increase in firing rate centered around 200 milliseconds (ms). In (C), the average of the responses across all experiments is plotted, with the black line generated from data from (B) and the gray line from data from (A). Note the separation in the two lines at about 100 ms after stimulus onset, indicating a difference in neural activity due to attention. Reprinted from J. H. Reynolds, T. Paternak, and R. Desimone (2000) "Attention Increases Sensitivity of V4 Neurons." Neuron 26: 703–14 with permission from Elsevier.

To characterize sensory neurons, neuroscientists use the concept of a *receptive* field. For visual neurons, the receptive field is that region of the retina where stimulation by light causes the neuron to change its firing rate relative to baseline. A neuron's receptive field can also be defined as an area of external space, in which a stimulus causes changes to the neuron's firing rate. Properly understood, this latter definition is equivalent to the definition that appeals to the retina, since the relevant



Figure 2.3 Location of two visual receptive fields. In the figure, the gray rectangle can be understood as the screen on which stimuli are projected. The eyes are at a distance of 57cm from the screen and fixated at the point where the two lines from the eyes converge. At that distance, one visual degree spans one centimeter. The visual receptive field of a neuron in primary visual area V1 is small, perhaps 0.5–2 visual degrees near the fovea. Receptive fields increase in size as one goes up the visual hierarchy. By the time one reaches the inferotemporal (IT) cortex where we have representations of objects, the visual field of an IT neuron is nearly 40 visual degrees in size. The figure is reproduced from "Visual Receptive Fields" at http://www.scholarpedia.org/ article/Receptive_field with kind permission from Jose-Manuel Alonso. The presence of a preferred stimulus within the receptive field drives the neuron to respond.

area of space projects to a specific region of the retina (see Figure 2.3).⁵ When I speak later of attending within the receptive field, I will use the second definition.

Furthermore, many early visual areas, such as primary visual area V1, are retinotopic in that neurons adjacent to each other have receptive fields that are adjacent to each other. As many studies on attention involve humans or macaque monkeys, the following is a map of some salient visual areas in both:



Figure 2.4a Human Visual System showing the four lobes of the brain: frontal, parietal, temporal, and occipital. Relevant cortical visual areas are as indicated. There are two cortical visual streams of note: the dorsal stream that extends into the parietal lobe (with distinction between a dorsal-dorsal and the ventral-dorsal pathway; extending into SPL and IPL, respectively); and a ventral stream that extends into the temporal lobe. MT (middle temporal area; also known as V5); MST (medial superior temporal area); IPL (inferior parietal lobule); SPL (superior parietal lobule); IT (inferior temporal cortex). Also included are visual areas V1–V6.



Figure 2.4b The macaque (Rhesus) monkey visual system along with areas related to eye movements and their connectivity: FEF (Frontal eye field); LIP (lateral intraparietal area); MST (medial superior temporal area); MT (middle temporal area). Reprinted from K. G. Thompson and N. P. Bichot (2005) "A visual salience map in the primate frontal eye field." Progress in Brain Research 147: 248–362 with permission from Elsevier.

To take stock, there are neurons in the visual system that have spatial receptive fields such that stimulation by light from an object in the external world can drive the neurons to generate spikes at a certain firing rate and thus to carry information about the stimulus. The questions then are this: What happens to neural activity during attention? What do these changes in activity have to do with the function of attention?

2.4 Attention-driven changes in activity in single neurons

There are a variety of measured effects on neural activity when attention is deployed, but let us focus on two: (1) changes in gain; and (2) receptive field remapping.⁶ We shall focus on visual neurons.

2.4.1 Attention and gain modulation

The intuitive idea behind change of gain is the boosting of a signal over background noise. Imagine standing in a field at night and dimly seeing an animal moving in front of you. Given the darkness, it is very difficult to identify the object, to separate it from the similarly dim background. Now imagine shining a flashlight on the animal, illuminating it. In this way, you boost the "signal" of the animal to the visual system, making it contrast more strongly against the background of the field. You can now see clearly that it is a fox. Might attention function like the flashlight, functioning also to boost a visual signal?

Many visual neurons have a preferred stimulus that drives them to generate a high firing rate. For example, visual neurons may prefer lines of specific orientations, motion in particular directions, or even certain objects like a fox. One can identify a neuron's preference by placing different types of stimuli in the neuron's receptive field and then recording the neuron's firing rate. Let's connect the example of the spotlight on the fox to the response of an actual neuron to visual contrast. This neuron responds to alternating black and white lines, i.e., high contrast. Furthermore, this neuron also prefers lines at specific orientations as shown in Figure 2.5. This response curve was generated by presenting the neuron with high contrast lines at different orientations and then plotting its firing rate for each orientation. It would be natural to think that neurons are only "on" for stimuli that they like and "off" for stimuli that they do not, but the figure clearly shows that neural response isn't simply a matter of being on or off. Rather, a neuron's firing rate is sensitive to variations in the stimulus with which it is presented, in this case variations in orientation. For this neuron, the preferred stimulus is one with lines oriented at 105 degrees (this is at the maximum of neural activity).



Figure 2.5 Multiplicative Gain: A plot showing the activity of a single neuron in visual area V4 to a visual stimulus. A "line" (actually a Gabor patch or contrast gradient) was presented at various orientations in the neuron's receptive field. The open circles show the activity to the stimulus when attention was directed outside the neuron's receptive field. The preferred stimulus is a line oriented at about 105 degrees. The black squares show the neuron's activity when attention was directed within the neuron's receptive field. Attention gain modulates response to the stimulus, amplifying its signal (compare the black and white points for each orientation). The dashed line shows the neuron's baseline activity in the two conditions in the absence of a stimulus in the receptive field. Reprinted from C. J. McAdams and J. H. R. Maunsell (1999) "Effects of Attention on the Reliability of Individual Neurons in Monkey Visual Cortex", Neuron 23: 765–73 with permission from Elsevier.

While theorists identify attention with selection, there are many kinds of selection that don't count as attention. Specifically, neurons can be selective without themselves paying attention. The neuron just discussed provides an example of this. It is clearly selective in preferring a specific orientation, but it would be odd to say that the neuron is thereby paying attention to the stimulus. Indeed, the lower curve in Figure 2.5 is measured in the *absence* of attention. What does this mean? In an experiment, one can manipulate where the animal directs attention by how one designs the experimental task. If one makes the lines in the previous example task *irrelevant*, but requires that the animal focus on a different part of space outside of the receptive field of the neuron, then the animal will not attend to the lines or, indeed, to the receptive field. This "minus-attention" condition is how one generates the first curve. What happens when the lines are *task relevant*, so that the animal now attends to them? One can again plot the neuron's response to the lines under this "plus-attention" condition and

the results are striking (see Figure 2.5, top curve). It is as if each point in the minus-attention curve is multiplied by some constant factor to generate the corresponding response in the plus-attention condition. The firing rate increases when attention is directed within the neuron's receptive field, relative to when attention is not so directed. This is understood as multiplicative or *response gain*: attention to a stimulus within a neuron's receptive field boosts the neuron's firing rate.

There is also a second type of change in gain, what is called contrast gain, as shown in Figure 2.6:



Figure 2.6 "Contrast-gain model of spatial attention. Spatial attention directed from outside the RF (light curve) to inside the RF (dark curve) shifts the contrast response function of the model neuron leftward. Responses are in arbitrary units." Reprinted from Geoffrey M Boynton, (2005) "Attention and visual perception" Current Opinion in Neurobiology. 15: 465–469 with permission of Elsevier.

The right curve shows the response of a neuron sensitive to contrast. When attention is deployed to stimuli in the neuron's receptive field, there is a leftward shift of the response function in the plus-attention condition. Let us simplify matters by noting that this shift in response reflects a greater sensitivity to contrast when attention is deployed. To anthropomorphize the neuron, it is as if the neuron can "see" lower contrast as if it were higher contrast. At lower levels of contrast, the neural response is more vigorous (remember the fox in the dark).

Why should these effects on single neurons be useful in thinking about attention as a selective process? Noudoost et al. (2010) write: "The objective of attention can be viewed as increasing the signal-to-noise ratio (SNR) of the readout from sub-populations of neurons encoding the selected representation" (p. 183). Think of signals as carrying information (in

Shannon's sense, see Appendix A), and noise as interference. Effective transmission of information, then, goes with an SNR ratio that is high, and one way to achieve this is to boost the signal relative to noise by changing the gain of appropriate neural signals. The intuitive idea is that when attention selects some stimulus, it amplifies representations of that stimulus (signal) over other representations (noise). Perhaps this is the neural correlate of a spotlight. That said, it is important to emphasize the limits of the spotlight metaphor. It is not as if attention just locally amplifies a signal, as if there was a spotlight on specific neurons in the brain. Rather, attention appears to have varied effects on neural activity, including activation and suppression (Datta and DeYoe 2009), and, even at the level of behavioral response, many psychologists long ago dropped the spotlight metaphor (Eriksen and St James 1986). The gesture at the spotlight metaphor as echoed in gain modulation is only to indicate some suggestive if inchoate connections between psychology and neuroscience. A less contentious point is that if there is anything to spotlighting, it has an echo in gain modulation.

2.4.2 Receptive field remapping

Receptive field remapping involves selectivity within the receptive field and, in a sense, a form of neuronal filtering. Let's first begin with the following question: Placing a preferred stimulus in a neuron's receptive field drives a robust neural response, but what happens to neural activity when there are two stimuli in the receptive field? Take a case where a neuron is presented with its most and least preferred stimulus, i.e., a stimulus that generates the maximum firing rate and one that generates the minimum firing rate relative to baseline. What does the neuron's response look like? You might expect the neuron to respond to the stimuli as if they were independent of each other and, thus, the total firing rate to be the sum of the firing rate to each stimulus presented individually. This is not what is seen. In visual areas V2 and V4 in minus-attention conditions, the firing rate was observed to be the weighted average of the two stimuli (Reynolds, Chelazzi, and Desimone 1999). This averaging is referred to as divisive normalization.

The averaging effect is seen in minus-attention conditions, i.e. when attention is directed outside the neuron's receptive field. One can now design an experiment where the animal has to direct attention to only one of two stimuli in the neuron's receptive field, either the preferred or unpreferred stimulus. Again, this is done with a well-defined task where only one of the two stimuli is task relevant. Moran and Desimone (1985) did this, recording from cells in area V4 and found that firing rate corresponded to the stimulus that was the target of attention as if the second stimulus in the receptive field was ignored. In other words, the response was no longer averaged but was seemingly driven only by the attended stimulus. To put it another way, if the animal was attending to the preferred stimulus, then the firing rate was as if that stimulus was the only one in the receptive field.

Moran and Desimone noted that "when attention is directed to one of two stimuli in the receptive field of a V4 cell, the effect of the unattended stimulus is attenuated, almost as if the receptive field has contracted around the attended stimulus" (p. 783). That is, it is almost as if the receptive field remaps, getting smaller so that unattended stimuli are effectively shut out. Echoing Broadbent's ideas, Moran and Desimone also speak of the effect as "filtering". One can thus rationalize receptive field remapping in terms of selective filtering and connect a basic feature observed at the hardware neural level with more abstract descriptions of attention provided by psychologists. This suggests a tantalizing way of bridging Marr's levels, namely the more computational mechanisms of the algorithmic level with the biological mechanisms of the implementation level. The bridge is through a common functional description: filtering. Nevertheless, a strong word of caution: these links are merely suggestive and the task of bridging the neuroscience and psychology of attention will be a complicated matter.

2.5 Neural-based theories of attention

I have thus far focused on individual neurons that echo the metaphors of attention as a spotlight and as a filter, but the psychological capacity of attention is implemented in the activity of populations of neurons and, indeed, of whole brain regions. This section surveys two influential mechanisms that look to the interaction between neurons: biased competition and divisive normalization.

2.5.1 Biased competition

Robert Desimone and John Duncan (1995) presented the concept of biased competition as central to understanding attention: "the model we develop is that attention is an emergent property of many neural mechanisms working to resolve competition for visual processing and control of behavior" (p. 194).

Strikingly, Duncan and Desimone do not identify attention with the underlying neural mechanisms, but with an "emergent" property. For current purposes, the proposal is that the mechanisms of biased competition realize attention. What then is biased competition?

Competition is connected to resource limitations, echoing concerns familiar from the psychology of attention. Focusing on the visual system, Duncan and Desimone (1995) suggested that the receptive field is a limited resource:

receptive fields can be viewed as a critical visual processing resource, for which objects in the visual field must compete ... If one were to add ever more independent objects to a V4 or IT receptive field, the information available about any one of them would certainly decrease. If, for example, a color-sensitive IT neuron were to integrate wavelength over its large receptive field, one might not be able to tell from that cell alone if a given level of response was due to, say, one red object or two yellow ones or three green ones at different locations in the field.

(197)

Desimone and Duncan identify the receptive field as a limited resource, but what they likely mean as the resource is the neural response of that neuron: namely, firing rate. Given biophysical limits, there is a limit to the number of spikes per unit time that can be generated. When two stimuli, X and Y, are presented in the neuron's receptive field, the neural response (firing rate) is observed to be the weighted average of the individual responses. This averaging is the manifestation of competition between the stimuli for the neuron's limited capacity for spikes. To the extent that a stimulus can win the competition for spikes, the sign of winning is that a stimulus, say X, gets the neuron to respond as if only X is in the receptive field while the loser, Y, gets ignored from the perspective of neuronal response. It is as if the receptive field has shrunk around X and excludes Y.

The crucial idea in respect of attention is that competition is biased. One stimulus is favored over another. To see how this biasing might work, it will help to take a provisional stand on these questions. In Chapter 1, I suggested that top-down attention involves goals influencing selection. A natural idea, then, is that what biases competition are goals, the tasks one aims to execute. These goals render some stimuli more relevant than others, and as a result of one's goals, the competition is biased in favor of goal-relevant
stimuli. But once there is goal-directed biasing of this sort, attention "emerges" or is "realized" in competition when the winner takes all of the spiking resources. Note that this conception of biasing is a way of connecting tasks to attention: the goals that the subject aims to achieve are tied to tasks that work towards that achievement.

Here is a speculative extension of the basic picture. Consider the visual system's being presented with stimuli X and Y, say two lines of different orientation. If the subject attends to X, then at appropriate stages in the visual processing hierarchy, receptive fields might remap in favor of X over Y due to biased competition, so that it is as if only X is present in the visual field. X has won the competition over Y. Let this process play out throughout the visual hierarchy such that at the end, for the sake of argument, all resources has been apportioned to the winner X. So, the system is in a state of having selected X rather than Y at the level of neural response. Thus, the resolution of competition for spikes can be seen in this receptive field remapping across relevant parts of the visual system, yielding the selection of X over Y for further processing. Perhaps this is how attention as selection of X for task emerges at the subject level, implemented by biased competition (this is an open empirical question). The speculative idea is a link between a schematic mechanism, resolution of competition for spikes and consequent remapping of receptive fields, and the selecting of a stimulus for further processing such as performing a task.

The expectation, then, is that large-scale neural activity will be apportioned to the winning stimulus and away from the losing stimulus in the competition for neural activity. Some experiments using imaging such as functional magnetic resonance imaging (fMRI) provide a snapshot of brain activity that is consistent with the predictions of biased competition. fMRI measures metabolic changes in the brain, specifically changes in blood flow through the blood oxygen-level dependent (BOLD) signal. The assumption is that such metabolic changes reflect changes in neural activity as such activity consumes energy. This requires that the relevant brain areas receive additional oxygenation. What is observed is that attention to X rather than Y activates regions devoted to processing X at the expense of those regions devoted to processing Y. For example, the observed BOLD signal can increase relative to baseline in regions of the brain that respond to X, while it decreases in regions that respond to Y (see Kastner and Ungerleider 2001 for discussion of some relevant results).⁷ In the Marrian framework, some forms of attention involve selection of information for task. This selection, described functionally, might be implemented by biased

competition that leads to receptive field remapping and the filtering of information.

2.5.2 Divisive normalization

As we have seen, neurons often show divisive normalization, seen in a weighted average response when multiple stimuli are placed within their receptive fields. In this section, I briefly sketch a link between attention and divisive normalization. Currently, this is an area of active research in neuroscience, and investigation is in its early stages. The current hypothesis is that the mechanism for attention acts through the mechanism that yields divisive normalization. The hope is that the appeal to divisive normalization might explain the neuronal effects of attention. If so, then the mechanism uncovered might be a fundamental neural mechanism of attention. To get a sense of how divisive normalization comes about, take a neuron, N, which exhibits normalization, an averaging of its response to multiple stimuli. It is postulated that N's response to multiple stimuli in its receptive field is the result of its interaction with other neurons that also respond to the stimuli and suppress N's response. Reynolds and Heeger (2009) speak of this suppressive population of neurons as N's normalization pool. Thus, the implicated mechanism yielding normalization can be modeled as dividing N's response by the total activity of its normalization pool (for recent models, see John H. Reynolds and Heeger 2009; Lee and Maunsell 2009; Boynton 2009). N responds to a stimulus in a way determined by its "listening" to other neurons that are also responding to that stimulus.

Visual neurons have a preferred stimulus to which they respond by firing strongly. A nonmathematical way to think about normalization is as a neuron's answer to the following question: How similar is the set of stimuli in my receptive field to my preferred stimulus? One can read the strength of the neuron's response as a report of stimulus similarity to the preferred stimulus, e.g., in respect of orientation, direction of motion, or some other feature. Thus, when the preferred stimulus is presented, the neuron fires strongly (the stimulus is of my preferred type); when the least preferred stimulus is presented, the neuron fires weakly (stimulus is not very similar to preferred); and when multiple stimuli are presented, then the response is some weighted average (the set of stimuli is somewhere in between preferred and non-preferred). The appearance of the weighted average when multiple stimuli are present is a neural signature of normalization.

Why is normalization important to attention? Note that not all visual neurons exhibit normalization (Lee and Maunsell 2009; Ni, Ray, and Maunsell 2012). For those that don't, one can think of their response as an answer to a stricter question: Is my preferred stimulus in my receptive field? Where the preferred stimulus is present, the neuron fires strongly, even if there are non-preferred stimuli present as well. Thus, a non-normalization neuron's response to multiple stimuli is not the weighted average of its response to each stimulus individually. Indeed, non-normalization neurons do not show the effects of attentional modulation that normalization neurons do. There is some data that suggests that the neural mechanism that implements normalization is also the one through which attentional modulation works. Specifically, attentional effects such as receptive field remapping are seen in (i.e., are correlated with) neurons that show normalization and are not seen in neurons that do not show normalization (Lee and Maunsell 2009, fig. 6). The suggestion, then, is that the effects of attention on neural activity might be mediated by the mechanisms of normalization.⁸ Indeed, divisive normalization has been used to explain the neural effects of attention noted earlier, such as changes in gain and receptive field remapping (again, see John H. Reynolds and Heeger 2009; Lee and Maunsell 2009; Boynton 2009). It is worth pointing out, however, that normalization is not restricted to attention. Matteo Carandini and David Heeger (2012), for example, argue that normalization is a "canonical neural computation," one that is computed throughout the brain. Normalization, then, might be a fairly ubiquitous feature of neural computation, one through which attention acts to generate some of the effects on neural activity noted earlier.

Let us now set the discussion in light of the Marrian framework. Neuroscientists have identified some effects of attention on neural activity: gain modulation and receptive field remapping. This is activity at the implementation level, measured directly from neurons. There are also more abstract computations and mechanisms such as normalization and biased competition. These ideas point to psychological correlates in terms of filtering and spotlighting. Finally, there is an overarching conception of attention from the empirical sufficient condition for attention: selection for certain tasks. A task for empirical science will be to tie all these disparate elements together in a comprehensive theory of attention. Certainly, that task of tying these ideas together is a large one, but there are enough potential links to suggest optimism that a Marrian explanation of attention might be possible.

2.6 The Premotor Theory of Attention

Let us now consider one last, influential theory of spatial attention in vision, the Premotor Theory of Attention. In primates, work in neuroscience has suggested that the top-down network of attention involves several areas: the frontal eye field (FEF) in the frontal lobe that connects to the lateral interparietal sulcus (LIP) as well as to other visual areas, including area V4. FEF is also connected to a subcortical structure, the superior colliculus (SC), which is necessary for generating saccadic eye movements, the ballistic movements that occur 2–3 times per second (for a succinct, somewhat technical overview of the relevant circuitry for eye movements, see Munoz 2002). Indeed, stimulating FEF neurons with a sufficiently strong electrical current can generate saccadic movements.

FEF also seems to have a role to play in the mechanisms of covert attention, namely attention that is independent of actual eye movements. This is suggested by an experiment by Moore and Fallah (2001). When a monkey's head is fixed and it is maintaining fixation on a central point, stimulation of FEF neurons with a strong enough current generates a saccadic eye movement to another location. Analogous to the idea of a receptive field, the region of space to which the eye movement is generated after stimulating an FEF neuron can be understood as identifying the motor field for that neuron, the space to which an action is directed. Moore and Farah had monkeys perform a detection task: to release a held lever when a target changed in luminance. The target was not at the fixation point, but within the receptive field of a neuron from which recordings were being taken. Using stimulation of FEF neurons at lower currents that did not generate eye movement, they found that detection performance was enhanced when FEF stimulation occurred while a target was present in the neuron's motor field, again at the location to which the eye would have moved at higher stimulation currents. So, activating an FEF neuron whose motor field is at X facilitates target detection at X. While stimulation at lower currents did not induce a saccade, it led to attentionlike effects on neural activity. The tantalizing idea is that stimulation at FEF was sufficient for covert attention to X (one should be cautious since stimulation typically effects a large number of neurons; Histed, Bonin, and Reid 2009).

Further work suggests that electrical stimulation of FEF neurons with a motor field at location X also yields attentional effects in neurons in other visual areas whose receptive fields are at X. That is, when the motor field of

an FEF neuron overlaps with the spatial receptive field of a visual neuron, electrical stimulation of the former can lead to attentional effects in the activity of the latter. Moore and Armstrong (2003) found that stimulation of FEF neurons at levels insufficient to induce a saccadic eye movement did increase the gain of V4 neurons whose receptive fields overlapped the FEF motor fields. Thus, stimulation of FEF neurons led to changes in V4 neuronal activity of the sort observed when attention is directed into the receptive field. Interestingly, stimulation of FEF neurons whose motor fields did not coincide with the V4 neuron's receptive field led to suppression of the V4 response. To echo James' observation, the idea of attention being increased and withdrawn has an echo in differences in neuronal activity. These effects in V4 were observed about 40 milliseconds (ms) after stimulation at FEF, which is consistent with FEF being the source of the changes. This provides suggestive evidence that attentional effects in V4 can be mediated by prior activity in FEF. In general, manipulation of nodes in the top-down attentional network such as FEF might lead to attentional modulation. In light of these results, one might wonder whether the representations of eye movement in FEF play a role in directing spatial attention. Put another way, might the mechanisms of overt attention, i.e., mechanisms yielding eye movement, be closely tied to the mechanisms of covert attention? The Premotor Theory answers this question in the affirmative.

The Premotor Theory of attention holds that there is a tight connection between the circuitry responsible for generating eye movement and that generating covert spatial attention. An influential conception of attention is that it is a general "supramodal" capacity that is distinct from but influences perceptual and motor processing. In contrast, Raymond Klein (1980) suggested that

When attention to a particular location is desired, the observer prepares to make an eye movement to that location; the oculomotor readiness ... has the effect of enhancing processing in or from sensory pathways dealing with information from the target location.

(262)

Klein spoke of this as the oculomotor readiness hypothesis (OMRH). Later, Giacomo Rizzolatti (1994) and coworkers proposed the Premotor Theory, arguing that attention is not tied to a supramodal system separate from

sensorimotor circuits, but that it emerges from those circuits. Klein and Lawrence (2012) suggest that OMRH and the Premotor Theory "both propose that the process of preparing to move the eyes ... is the mechanism by which endogenous covert spatial attention is engaged" (15). As the Premotor Theory has been more widely discussed, I shall focus on that.

A central idea in the Premotor Theory is that of a spatial pragmatic map. Such maps are embodied in neurons that have both motor and visual receptive fields that spatially overlap. Again, a motor field of a neuron N is that region of space L such that a motor output towards L triggers the neuron to fire. Let the visual receptive field of N also be at L. Then the visual and motor receptive fields overlap. The neuron thus fires either when a stimulus is placed in L, activating its visual receptive field. One can then think of this neuron as coding both a motor and visual spatial map of L, i.e., a spatial pragmatic map.

Rizzolatti et al. (1994) note three central tenets of the Premotor Theory:

- 1. The mechanisms responsible for spatial attention are localized in the spatial pragmatic maps ...
- 2. Spatial attention is a consequence of a facilitation of neurons in the spatial pragmatic maps. This facilitation depends on the preparation to perform goal-directed, spatially coded movements.
- 3. Different spatial pragmatic maps become active according to the task requirements.

(239 - 40)

There is then a simple construal of the Premotor Theory. If spatial pragmatic maps do double duty as action and spatial representations, then when the subject intends to execute the represented action, this activates the action representation in the spatial pragmatic map. But since these representations are spatial pragmatic maps, the activation of the action representation of necessity activates the spatial representation. Assuming that this activation of spatial representations suffices for the inducement of spatial attention to the region in question, then appropriate action yields appropriate attention.⁹ This is a striking way to tie attention to task, for on the Premotor Theory, spatial attention to a location is realized in the activation of specific action representations to that location. That is, in selecting a location as a target for action, the subject thereby attends to that location. There are three possible formulations of the Premotor Theory: the identity, causal, and anatomic formulations. The strongest thesis is the identity formulation: the neural circuits for preparing eye movement to location L are just the neural circuits for visual spatial attention to L. The causal formulation is weaker in that it endorses only a causally sufficient condition: normal preparatory processing for eye movement to location L is sufficient to cause visual spatial attention to L. Finally, the anatomic formulation holds that a brain region contains circuitry for preparatory movement activity and for spatial attention.

Although the results from Moore and co-workers are consistent with the identity and causal formulations of the Premotor Theory, there does seem to be compelling data against the identity formulation. For example, Pouget et al (2009) have argued on anatomical grounds that the FEF neurons that influence V4, as in the Moore and Armstrong results, are distinct from the FEF neurons that play a role in eye movement. The relevant FEF neurons for both circuits occur in different layers of the cortex (the cortex has six layers). If this result holds up, then the FEF neurons subserving attention are anatomically distinct from those subserving eye movements. So, there is an anatomical distinction that undercuts the identity claim regarding motor-attentional circuitry. There seems to be growing consensus that the current body of evidence speaks against the identity formulation of the Premotor Theory (Smith and Schenk 2012). The causal formulation, however, is left open and further work is needed to decide whether it is correct. Thus, it might be that some of the mechanisms of spatial attention involve circuits tied to preparing eye movement. There is also reason to think that the anatomical formulation might be correct, specifically given the role of FEF in both the generation of eye-movement and the inducement of spatial attention. Still, to the extent that the anatomical Premotor Theory is correct, it is on its own mechanistically uninteresting, for it amounts to the claim that certain brain regions contain neurons involved in generating eye movement and neurons involved in generating spatial attention.

An interesting feature of the Premotor Theory is its denial of a supramodal account of spatial attention that emphasizes the attentional system as distinct from perceptual circuits. If the identity formulation of the Premotor Theory is false, then the supramodal theory of spatial attention remains on the table. In general, the question of whether the capacity for perceptual attention is implemented in supramodal circuits or in perceptual circuits remains an open empirical question.

2.7 Challenges to a neural mechanism-based metaphysics of attention

Is attention a brain process? Might it be a specific neural mechanism? We are ready to assess these questions. Any neural mechanism-based approach to the metaphysical question has two implications:

- Necessity: Where attention obtains, so does the relevant neural mechanism.
- Sufficiency: Where the relevant neural mechanism obtains, so does attention.

A neural mechanism-based approach to the metaphysics of attention might be an identity theory: attention = neural mechanism N. There are other claims that also imply necessity and/or sufficiency, say that attention is constituted by or supervenes on N, but for ease of discussion, I'll stick to the identity formulation as the stalking horse. This won't matter save for expository purposes. The targets are necessity and sufficiency. Philosophers have to date largely given up on mental-physical identities, at least those that identify a type of mental state with a type of physical state (a type-identity theory). Historically, an influential argument against the identity theory was the possibility of the multiple realizability of the relevant mental state. Here's a quick gloss of the argument: pain cannot be identified with a neural process, for pain could be realized in non-neural processes (e.g., one might be able to build an artificial system that feels pain). The possibility of alternative realizations of pain refutes necessity: the presence of pain does not entail a specific neural process. So, the identity theory fails.

Psychologists have made similar arguments regarding attention. Some years ago, Alan Allport argued against a broader mechanism-based meta-physics using a version of multiple realizability (broader in that he did not limit himself to neural mechanisms). He wrote:

most contemporary theories of information processing in general, and selective attention in particular, view attention as some sort of causal mechanism. However, even a brief survey of the heterogeneity and functional separability of different components of spatial and nonspatial attentional control prompts the conclusion that, qua causal mechanism, *there can be no such thing as attention.* There is no *one* uniform

computational function, or mental operation (in general, no *one* causal mechanism), to which all so-called attentional phenomena can be attributed.

(1993, 203)¹⁰

Allport's argument can be simply stated:

- 1. "The evidence ... indicates that, in implementing visual-attentional selectivity, a number of qualitatively different mechanisms are involved" (202).
- 2. So, the evidence is that there is no unitary mechanism that is or constitutes attentional selection.

Allport's argument is an empirical one. He takes the support for the first premise to be derived from experimental work of the sort discussed previously. But can one draw Allport's conclusion?

There have been philosophical concerns raised about the multiple realizability argument and questions about whether there are cogent versions of it (for a critical discussion, see Shapiro 2004; Shapiro 2000). In fact, I do not think that there is yet a cogent version against the identification of attention with a neural mechanism. To successfully run the multiple realizability argument against the claim that attention is a brain process, one must have: (A) a handle on the realized capacity, namely attention; and (B) a concrete description of two neural mechanisms, N and M, such that N and M are distinct in the requisite way and are plausible realizers of attention. (B), however, has not been fulfilled. In discussing the neuroscience of attention, I have identified several different neural effects that are observed under conditions of attention. But while these components are suggestive of what a complete mechanism for attention might look like, detailed proposals that delineate the complete mechanism that implements attention are needed. Having parts of a mechanism falls short of assembling them into concrete wholes, yet the multiple realizability argument requires two concrete and distinct mechanisms. Given the diversity of components that likely contribute to the total realization of attention, one can still wonder whether there will be a single mechanism type for attention. I share Allport's skepticism, but skepticism does not an argument make. Without concrete mechanistic proposals, the required premise needed for Allport's conclusion is not in hand. Neuroscientists need time to uncover these mechanisms.

To provide these mechanisms, however, neuroscientists must rely on the empirical sufficient condition, and this suggests that the connection between attention and task is fundamental to the empirical understanding of attention. After all, to run the multiple realizability argument, experimentalists must have a handle on attention and on its neural realization. Yet it is the empirical sufficient condition that provides a handle on both. One can get at the issue with this question: Why think that a change in neural response like multiplicative gain has anything to do with attention? Is there something inherent in that response that entails attention? Prima facie, no such entailment holds. After all, multiplicative gain is just a form of signal enhancement that can be implemented in non-attentional systems. Think of an amplifier that boosts the playback of notes of a recorded musical scale in such a way as to mimic the amplification of the sound to look like the top curve in Figure 2.5 where each note is amplified in a distinct way, and one note is the "preferred" stimulus, being played back the loudest. So, why tie gain modulation to attention? The reason is that neuroscientists observe gain modulation in the context of experimental subjects performing a well-defined task that deploys attention. In many of the experiments discussed previously, the experimental animals were required to perform tasks that involved their selecting targets inside or outside the recorded neuron's receptive field. Measurements were then made of neural response to targets within the neuron's receptive field. It is in light of the empirical sufficient condition that one can interpret the neural response as "attentional" as opposed to something else. Selection for task allows experimenters to infer that the previously described neural effects are linked to attention. Accordingly, the construction of any argument for or against a neural-mechanism based account of attention relies on the invocation of the empirical sufficient condition. This suggests, to me at least, that the connection to task is fundamental: it allows neuroscientists to keep a handle on attention in experimental contexts and to tie attention to specific neural effects. To construct a multiple realizability argument, one will need to rely on the appeal to tasks.

Multiple realizability raises questions about necessity, but I want to conclude by considering a different way one might try to unify the disparate neural effects previously noted, namely by appeal to divisive normalization. This will raise questions about sufficiency and further highlight the central role of the empirical sufficient condition. Christopher Mole (2011) has provided an argument against sufficiency as follows, here focusing on feature binding in Treisman's Feature Integration Theory:

- 1. If attention is a process, then, for all events x and y, if x and y instantiate the same process, then if either one of them is an instance of attention, the other is too.
- 2. There are some events that are instantiations of the feature-binding process and that are instances of attention.
- 3. There are some events that are instantiations of the same feature-binding process and that are not instances of attention.

Therefore

4. Attention is not a process.

(adapted from Mole, 2011, 46)¹¹

One can replace "process" with "neural mechanism" and in place of featurebinding substitute any appropriate mechanism. Mole does not consider neural mechanisms explicitly, but here is one possibility. Divisive normalization has been used to model both types of gain modulation (contrast gain and multiplicative gain). Indeed, some versions of divisive normalization are adapted from models of biased competition (Reynolds, Chelazzi, and Desimone 1999), and so it might be the case that divisive normalization can account for the effects associated with gain modulation and biased competition. This yields something of a hierarchy: signal modulation (spotlighting) and receptive field remapping (filtering) as underwritten by divisive normalization. The hypothesis then is that divisive normalization provides the fundamental mechanism for attention, explaining all attentional effects.

Is the deployment of divisive normalization sufficient for attention? The answer seems to be no. Some theorists who have invoked divisive normalization to explain the neural consequences of attention also claim that it is a *canonical* neural computation, meaning that it underwrites many other non-attentional phenomena (Carandini and Heeger 2012). That is what makes it "canonical". Accordingly, one can run Mole's argument against divisive normalization in respect of sufficiency by inserting this point at premise (3): there are instances of divisive normalization that do not entail attention. Attention then cannot be identified with divisive normalization. A natural rejoinder at this point would be that there are specific deployments of divisive normalization that yield attention, and the neural mechanism-based metaphysics will appeal only to these deployments. Fair enough, but what is the additional ingredient that renders certain deployments attentional? One possibility is that neuroscientists will find a neural property that renders certain instances of divisive normalization as uniquely attentional. Perhaps, but how would one know what that property is? How does it signal for attention? There is, again, no way out but to appeal to the empirical sufficient condition: it will be just those deployments that are observed during performance of specific tasks. Specifically, the claim is that divisive normalization specifically gives rise to attention when it serves selection for task.

This recalls the Marrian point that a theorist's handle on how neural activity is related to psychological phenomena is dependent on understanding the purpose or function of the psychological phenomena. In the case of attention, the core assumption across experimental approaches is selection for task. The natural suggestion, then, is that the uniform feature that points to attention for each relevant neural phenomenon just is selection for task. Attempts to assess a neural mechanism-based metaphysics of attention seem to always revert to the empirical sufficient condition. Might it, then, just be that tasks are what is in fact central to answering the metaphysical question regarding attention?

2.8 Summary

Our short tour of a vast empirical literature on attention in psychology and in neuroscience is now complete. This has included discussion of a number of experimental paradigms used to study attention, a set of prominent theories of attention including metaphorical characterizations, neural mechanisms, and more abstract computational processes like feature integration, divisive normalization and biased competition. In covering this material, we saw some interesting connections between different levels of analysis: Broadbent's filtering conception as echoed by receptive field remapping in neurons; the spotlight of attention as echoed by gain modulation in neurons; the empirical sufficient condition explicated in terms of a premotor account of spatial attention, and so on. These results have provided many detailed answers to our basic question about the mechanisms of attention, how attention works and how it is implemented. A challenge for the cognitive science of attention will be to bridge these links between neuroscience and psychology, but more broadly to provide some order and unity to the ever-growing empirical literature of attention. Given the skepticism about attention noted in the Introduction, the field needs unifying principles to tie together psychological work on behavior during attention, computational modeling of such data, and the fundamental mechanisms of attention as implemented in the brain. It is hard to see how genuine progress will be achieved without unifying principles. After all, questions and theories of attention will only be as clear as the central concept used to state them, but doubts about attention grow.

The fundamental lesson of the last two chapters then is this: one principle for unification has emerged, namely the idea of selection for task. The empirical sufficient condition provides a partial answer to the function question about attention. This condition is not idiosyncratic or the product of some decision by fiat by a subset of researchers. I have argued that it is built into experimental practice both in psychology and neuroscience. It is the only way to get a handle on attention, and it guides interpretation of the data. Since this condition is shared by theorists of attention, it provides a common anchor to ground subsequent research. That there is such a shared principle might be enough to do quite a bit of work in the cognitive science of attention. That is something that should be explored in cognitive science: where skepticism about attention is raised as a substantive issue in responding to theorizing or experiments on attention, one should set forward the empirical sufficient condition as a response. The next question, however, is philosophical: Might the selection for task condition provide a plausible answer to the metaphysical question? That is, might attention in some way just be selection for task or, more generally, for action?

Suggested reading

Marr's seminal discussion of his explanatory scheme is given in (1982, chap. 1). A summary of the neuronal effects of visual attention is given in Bundesen and Habekost (2008). Wright and Ward (2008, chap. 7) and Armstrong (2011) give a detailed discussion of the physiology of orienting attention. Ruff (2011) provides a discussion of biased competition. For a critical discussion of the premotor theory, see Smith and Schenk (2012).

Notes

- 1 A paper that got many philosophers to think more seriously about mechanisms is Machamer, Darden, and Craver (2000).
- 2 Specifically, Marr characterizes the computational theory in light of the following questions: "What is the goal of the computation, why is it appropriate, and what is the logic of the strategy by which it can be carried out?" On the representation and algorithmic level, the relevant question is:

"How can this computational theory be implemented?" In particular, "What is the representation for the input and output, and what is the algorithm for the transformation?" On the level of hardware implementation, the question is: "How can the representation and algorithm be realized physically?" These characterize "the three levels at which any machine carrying out an information processing task must be understood" (1980, p. 35).

- 3 Talk of a function of serving an experimental task might sound odd since we might think of functions in a teleological sense. The idea here is meant in more of a computational sense, but one development to be examined in the next chapter takes attention to be a function of selection for "task" in a much broader sense, say action in general. There, perhaps, we might have a psychological capacity that was selected by evolution.
- 4 Both MEG and EEG provide more fine-grained temporal information about brain activity. EEG is less expensive but suffers from distortions to the signal from the skull. MEG avoids that distortion, but is much more expensive to run. I'm grateful to Linda Moya for information on these forms of imaging.
- 5 Visual receptive fields can have a substructure: different neural responses are elicted depending on which part of the field is stimulated. For example, retinal ganglion cells have a *center-surround* structure where activation of the center can activate or suppress firing rate, while the opposite is found in the surround. We shall set this point aside. We focus on *spatial* receptive fields, but receptive fields need not be spatial. They can also be tied to a class of stimuli appropriate to driving the response of a neuron. Accordingly, olfactory neurons have a receptive field, but these are identified in terms of the class of chemicals to which they respond, and not in terms of some area of space (Koulakov et al. 2011). Thus, the notion of a receptive field is not exclusively a spatial one.
- 6 Other reported effects include the increase in the base-line activity of neurons (their base-line firing rate is greater under attention conditions (Luck et al. 1997)) and the synchronization of activity between neurons that some have claimed to be a fundamental mechanism of attention (Womelsdorf and Fries 2007). These other effects merit additional discussion but we do not have the space here to do so. Synchrony especially has gotten much recent attention, but while the basic idea can be glossed in terms of neural activity being jointly "in phase" (think of two sine waves being in phase), there are challenges to defining synchrony precisely, challenges as to how to measure it, and challenges in showing that it is

causally relevant and not merely an epiphenomenon. Given these complexities, we have to set this hot issue in current neuroscience aside. I am especially grateful to discussion by the CNBC Attention Group on this issue, and to Rob Kass and Linda Moya for presenting these ideas to the group, and to Steve Chase, Byron Yu, Raj Ghandi and Matt Smith for contributions.

- 7 fMRI reveals global activity of large regions of the brain. A "voxel" (from "volumetric pixel"), the basic unit of analysis in fMRI, typically covers the activity of millions of neurons.
- 8 The contrast between neurons that exhibit normalization and those that do not is perhaps a bit stark, and it is likely that there is a continuum.
- 9 One might wonder if the theory is limited to visual spatial attention. In principle, the theory could be extended to other modalities and targets of attention. First, it is not just locations that are of relevance to actions, but also objects and features: we choose the red rather than green apple, or the grape rather than pear. Perhaps there are motor fields tied to features and objects. Might there then be feature or object pragmatic maps? Second, many of the other modalities seem to have motor correlates with attention: a loud bang or a soft sound behind the wall induces a head movement to aid auditory localization or a prick on one's arm induces one to reach for the site of pain. Might there then be auditory or tactile spatial pragmatic maps? If the Premotor Theory is to be more than a theory of visuospatial attention, these questions have to be answered.
- 10 In his discussion, Allport highlighted a diverse number of processes that he judged to undercut the identification of attention with a specific causal mechanism, including some of the ones that we have noted (202–3): spatially selective enhancement of neuronal response, raised threshold of motor responsiveness, enhanced selectivity of tuning, increased local competition, task related suppression of prepotent response tendencies, maintenance of executive working memory, temporal and conditional sequencing of cognitive operations, and so on.
- 11 Mole argues against sufficiency by drawing on the phenomenon of *hemi-spatial neglect* where subjects are thought to be unable to attend to (typically) the left side of visual space. There are striking cases where neglect patients show some responsiveness in their behavior to information on the left that they cannot report. In a famous study with patient P.S., John Marshall and Peter Halligan (1988) presented a neglect patient with pictures of two houses. The central difference between the two houses was that in one, a fire can be seen emanating from the left side, the side that

is neglected. P.S. reported no visual difference between the two houses, and yet when prompted to choose between them, she often chose the one that was not burning. Similar results have been obtained with other patients where they evince sensitivity to information from the neglected portion of space. The subject cannot attend to the relevant region of the picture, yet there is presumably some feature binding of the elements of that region to allow for a "sense" of the presence of a fire. Thus, neglect presents a case where attention is by hypothesis absent but some form of feature binding is present. If so, attention can be dissociated from feature binding.

3

AGENCY AND THE METAPHYSICS OF ATTENTION

3.1 Introduction

In the last two chapters, selection for task has constantly emerged as a crucial condition and organizing principle. Moreover, given that the condition is a background assumption of experimentalists working on attention, as I argued, it provides an antidote to the despair of finding an answer to the metaphysical question that has afflicted some theorists of attention. This chapter pursues the optimistic path that selection for task, and indeed selection for action, can provide an answer to the metaphysical question.

I begin in Section 3.2 with Alan Allport and Odmar Neumann's foundational presentation of the selection for action view of attention. Section 3.3 then examines the basic case they consider, what I have elsewhere dubbed the Many-Many Problem. I use this Problem to motivate expanding the empirical sufficient condition, which ties attention to task, to a more general condition that ties attention to mental and bodily action. In Section 3.4, the empirical sufficient condition is expanded by stages to the general condition (S): if S selects X for action, then S attends to X. Section 3.5 considers situations where selection for behavior does not count as attention and explains the difference in terms of subpersonal- versus personal-level states. Section 3.6 then argues that action implies attention. In Section 3.7, I defend (N): If S attends to X, then S selects X for action. The status of (N) is that it explains how attention arises from mental states. I then defend (N) against two putative counterexamples: (a) attentional capture and (b) attention as vigilance. The selection for action account is then stated and interpreted in Section 3.8. I suggest that it is the best working hypothesis for explaining what attention is. Finally, in Section 3.9, I discuss Christopher Mole's cognitive unison account of acting attentively. I propose that cognitive unison conjoined with selection for action provides a more complete account of attention.

3.2 Selection for action: original psychological accounts

Theories of attention now and in the past have emphasized capacity limits as necessitating attention, leading to a conception of attention as selection that is required to avoid information overload. In arguing for attention as selection for action, Odmar Neuman and Alan Allport dissented from this consensus that attention is necessitated by capacity limitations. This is not to say that they denied the existence of capacity limits. Neumann (1987) emphasized two notions of capacity limitation. First, capacity limitations are just the observed limits to performance during standard attentional tasks, such as dichotic listening. This notion of limitation is uncontroversial and precisely what the theory of attention must explain. Second, Capacity limitation (which Neumann intentionally capitalized) refers to a theoretical construct that is invoked to explain empirically observed capacity limitations. According to most theorists, Capacity limitations necessitate selection processes so as to avoid information overload. In contrast, Neumann emphasized that "independent of all Capacity considerations, selection is evidently needed for the control of action. Organisms must constantly select what to do and how to do it" (374). Thus, the central problem that forces attention on the scene is

how to avoid the behavioral chaos that would result from an attempt to simultaneously perform all possible actions for which sufficient causes exist, i.e., that are in agreement with current motives, for which the required skills are available and that conform to the actual stimulus situation.

(374)

In other words, constraints from action rather than Capacity limitations necessitate attention.

For Neumann, two basic problems are raised by actions, understood as flexible, non-reflex behaviors: effector recruitment and parameter specification. Effector recruitment is connected to a concrete limitation, for an agent has only a limited number of effectors (body parts) to deploy. Moreover, deploying effectors involves basic skills that are applicable to different circumstances. To implement skills in concrete circumstances, an agent has to set specific parameters that allow their appropriate expression. In the case of perception-guided action, this involves in part drawing on relevant information from the environment to inform an appropriate response. Thus, a subject knows how to reach for objects, but to implement a reach for a specific object, the context will determine what is the appropriate reach, type of grip, trajectory, and grip force based on information about the object and the environment given the bodily effector selected. To deploy a recruited effector in specific circumstances, the precise parameters of the movement must be specified, though such specification can be a complex, temporally extended process that occurs throughout the movement. For example, information from vision about object location, shape, texture, weight, etc. will guide (set parameters for) a reach and grasp movement with one's right hand.

Contemporaneously, Alan Allport (1987) presented a theory of attention as selection for action. Here is a simple case he presents:

Many fruit are within reach, and clearly visible, yet for each individual reach of the hand, for each act of plucking, information about just one of them must govern the particular pattern and direction of movements. The disposition of the *other* apples, already encoded by the brain, must be in some way temporarily decoupled from the direct control of reaching, though it may of course still influence the action, for example as representing an obstacle to be reached around, not to be dislodged, and so on. The same necessity of selecting, in respect of a given class of action performed by a given effector system, just *one* among a number of physically available objects to act upon appears to be essentially universal ... Although the senses are capable of registering many different objects

together, effector systems are typically limited to carrying out just one action of a given kind at a time. Hence the biological necessity, and the theoretical importance, of selection-for-action.

(396-97)

Moreover, Allport notes that "the need for such a mechanism (of selective coupling and decoupling of perceptual and motor processes) arises directly from the many-many possible mappings between domains of sensory input and of motor output within the very highly parallel, distributed organization of the nervous system" (397).¹ This central but neglected insight provides the basis both for a philosophical theory of the nature of action (as developed by Wu 2011), and an account of the nature of attention as selection for action. As I develop the Allport and Neumann model, the emphasis will be on Allport's notion of coupling an input to output in action, where this coupling involves what Neumann spoke of as parameter specification and effector recruitment, namely, the use of the input to guide, control, and program an appropriate output. For example, relevant visual information is used to guide a specific train of reasoning. Coupling explains selection as for action.

3.3 The Many-Many Problem and selection for action

In this section, I explain Allport and Neumann's idea of selection for action within what I call the Many-Many Problem. To see why Allport and Neumann think that agency necessitates attention, consider a simple situation where one can move and manipulate an object, thereby performing an action. For example, there is a basketball before me, and I kick it with my dominant right foot. Easy enough, but the simplicity of this action obscures a structure that seems to underlie all mundane bodily actions. This structure is revealed by considering a slightly more complicated situation: not only is there a basketball, but also a soccer ball. Moreover, I can kick with either foot, so I have more than one available action: I can kick either ball with either foot. These options delineate a set of behavioral possibilities that includes four possible kicking actions. Let us restrict this behavior space to just the two targets and two effectors (my two feet) to keep things manageable. This will suffice to make all the points needed at the outset. One can then depict as follows the behavior space, the space of possible actions for a subject at a time and location, for the kicking example:



Figure 3.1 A simple behavior space with two inputs, a soccer ball and a basketball, and two potential outputs, kicking with the left foot or kicking with the right foot. There are four corresponding possible actions. Note that on the input side, there is a visual experience of that input. The agent visually experiences both balls.

This behavior space identifies a network of possible actions, given inputs and outputs available to the subject at a given time. The four couplings, represented by the arrows, identify the four possible actions. When the subject "takes one of the paths" depicted by the arrows, then that action is performed, and the arrow represents a causal process whereby the input informs production of the output to which it is coupled.

Notice, however, that I cannot perform all four actions simultaneously. This is generally true in cases of action. When one acts, the action implies selection of a path among multiple behavioral possibilities. What happens, then, when one acts? One traverses one of the available paths in behavior space, the implementation of a specific input-output mapping such as the perception-guided kick of a soccer ball with one's right foot. As Allport puts it, there is a selective coupling. If so, then actions are selective in being or entailing the subject's traversing one of the available paths in behavior space. So kicking the soccer ball with my right foot implies that this action occurs against the other possibilities available to me. I kick the soccer ball rather than the basketball, and with my right foot rather than with my left. Accordingly, the soccer ball becomes the target of my action, and to act on it, I must use appropriate visual information from it to guide my kicking response. This means that there has to be selection regarding relevant features of one ball, rather than the other, to guide movement. It is in the selection that one will find attention, namely selective attention to the ball and its features so as to inform the kick. This is the proposal to be unpacked.

The general idea is that bodily action entails traversing one path among many in a subject's behavior space. This space has a structure, a set of paths defined by the available targets (the inputs) and the available responses (the outputs). Each path corresponds to a potential action, a mapping of one among many inputs to one among many outputs (one can allow for multitasking in the sense of multiple paths that can be taken). The challenge for the agent is that this behavior space presents a Many-Many Problem: which path should be taken? That is, what should one do? To solve the Many-Many Problem is to traverse a path, one that ties a specific input to a specific output. In other words, the solution entails action, a response informed by the way the subject takes things to be (later, I will include one-many and many-one mappings as instances of the Many-Many Problem).

Two aspects of the Many-Many Problem should be emphasized: (1) the behavior space used to explicate the Problem is a psychological space; and (2) a crucial element is the notion of coupling. First, when one speaks of selection for action, "action" is being used in a special sense. In cognitive science, it often refers to physical movement. In philosophy, it is broader, referring both to mental and bodily actions, where these can involve additional properties such as responsibility, rationality, freedom, and intentionality. Actions are performed by agents, i.e., subjects capable of being in mental states.² Accordingly, the behavior space at issue is a psychological space: the inputs that structure the behavior space are items to which the agent stands in psychological relations, say appropriate perceptual, cognitive, or memory states. For example, the subject perceives or remembers the relevant input. The significance of this restriction will be discussed in Section 3.5.

The second point is to deemphasize the many-many structure of the initial presentation of the Many-Many Problem in favor of the notion of coupling. A behavior space with a many-many structure is a vivid way to illustrate why selection is necessary in many mundane actions. Never-theless, a many-many mapping is not essential. Rather, what is central is the need for input-output coupling. The critical moment in agency is whether an input is to be coupled to a behavioral output or not. A similar feature is found in a *one-many* behavior space, a situation where there is only one possible target of action, but many ways one can act on it. So, one might only have the soccer ball, but one can kick it, throw it, roll it and so on. Again, action requires selection, a specific input-output link. The same is true with a *many-one* behavior space where there are many targets although only one behavioral response. In fact, the point applies to a putative *one-one* behavior space so long as the action is something that need not be

done, so in effect there remain two options: to act or not to act (the behavior is not then a reflex, see Section 3.6). The point is that for all these "variants" of the Many-Many Problem, action entails the selective coupling of an input to an output. In what follows, I will continue to speak of a Many-Many Problem, though in doing so, I will always have in mind all the variants where solving the Problem involves input-output coupling. The Many-Many Problem could be renamed the Selection Problem.³ Elsewhere, I have argued that action just is solving the Many-Many Problem (Wu 2011b), but in the current context, the Problem serves to explicate the idea of selection for action via input-output coupling.

Bringing these two issues together, the crucial point is that the input state to coupling is a personal-level state where this state is attributed to a subject. This input state then initiates a process of coupling that can be quite complex. The role of the input state is, at a minimum, to set initial parameters in the parameter specification and effector recruitment needed to generate an appropriate action. For example, the input state can define the target for a reach or the spatial features needed to determine an appropriate grip. This specification contributes to the guidance of the output. The precise nature of the coupling process is, in the end, a matter for empirical investigation, but it is compatible with input-output coupling to solve the Many-Many Problem that coupling involves unconscious states or, indeed, states that do not involve the subject, such as certain neural states that realize coupling. This does not undercut the idea that certain psychological states function as the basis for input-output coupling, the starting point for generating action. The arrows in a behavioral space identify couplings that lead to action, an input-output mapping that amounts to an agent's response (output) in light of how the agent takes things to be (input).⁴ Selection for action is spelled out in terms of coupling that solves the Many-Many Problem posed by a behavior space that is also a psychological space.

To bring some distinctions explicated in Chapter 1 to bear on the current discussion, one can think of selection for action as sometimes constrained by intention. Where it is, then the relevant selection involves control; where it is not, then the relevant selection involves automaticity. The causal role of intention is conceived as aiding the solution to the Many-Many Problem. For example, that the subject intends to kick the soccer ball provides constraints to aid solving the Many-Many Problem, for the intention specifies the target and relevant response (for some further elaboration of this, see Wu 2008). For current discussion, the precise causal relation between

intentions and solving the Many-Many Problem can be left open in favor of the plausible idea that an agent's goals do influence the selectivity that is needed in producing action. In the kicking case, the fact that the agent intends to kick that soccer ball explains the agent's perceptual selectivity in respect of that ball. Accordingly, the agent's selectivity is controlled (later I shall say that the selectivity entails the agent's attending to the ball). Of course, many aspects of the agent's selectivity might be automatic in that they are not intended. As shown in Yarbus' experiment (see Chapter 1), subjects will move their eyes in goal-appropriate patterns even if they do not intend to make such specific patterns of movement. The implementation of those patterns is automatic. Similarly, the kicker might not only attend to the ball but to a specific location of the ball to make contact with the kicking foot, but that further selectivity of a location or a feature of the ball need not be intended and, hence, can be automatic. In the limiting case where intentions are not involved at all in input-output coupling, a process is fully automatic. This possibility points to attentional capture (see Section 3.7).

3.4 Selection for action as sufficient for attention

My goal in this section is to argue that the empirical sufficient condition,

 (S_{emp}) If subject S perceptually selects X for task T, then S perceptually attends to X,

can be expanded to the following conditional:

(S) If subject S selects X for some action A, then S attends to X.

In other words, I will motivate dropping the restriction to perception and to tasks in the antecedent, yielding a general sufficient condition for attention. Let us first move from tasks to bodily action. The relevant tasks at issue in S_{emp} , as discussed in Chapter 1, are actions like visual search, verbal shadowing, and reporting the presence of stimuli by some form of report (e.g., verbally or pressing a button). It seems unduly narrow, however, to restrict the notion of task in the sufficient condition to these specific experimental paradigms. After all, in performing these tasks, subjects just do the sorts of things they do all the time: they say things, look, listen, and produce responsive movements. While psychologists have focused on a

84 AGENCY AND THE METAPHYSICS OF ATTENTION

specific set of tasks in investigating attention, there is nothing special about the tasks that give psychologists special access to attention. Rather, such tasks are experimentally useful. Being tightly defined, they render data collection and interpretation easier. Given that the behavioral capacities that underwrite performance of experimental tasks are of the sort routinely performed in mundane actions, there is no principled reason to divide experimental tasks from mundane bodily actions such as kicking a ball. This suggests an expanded sufficient condition:

 (S_{emp1}) If S perceptually selects X for bodily actions, then S perceptually attends to X.

The leap from the original empirical sufficient condition to $S_{\rm emp1}$ is small, significant, and plausible. 5

Let us now provide reasons to drop the restriction to perceptual selection by noting that actions are often guided by non-perceptual mental states where this guidance involves a selection of X to inform action. A pervasive case is memory-guided action. Consider a simple case where you see several objects before you, but suddenly the light goes off. You were just about to reach for an object that you can no longer see in the dark. Nevertheless, you can remember where the object was when you last glimpsed it, and you use your memory of its spatial location to guide your reach to it. Indeed, you remember where the other objects were before the lights went out, so in using memory of the target object, you exemplify a selectivity in respect of what you remember that is aimed towards directing an action targeted at the object. This sort of case might be less than ordinary since it depends on sudden power failures, but the selective reliance on memory that it illustrates is mundane. Think of memorizing a shopping list and working through the list as you shop.

When psychologists speak of memory, they make a division between long-term and short-term memory. The relevant memory in guiding action is a form of short-term memory called working memory, or colloquially, memory for work.⁶ Long-term memories are not explicitly held in mind unless one recalls them, e.g., to serve action. This then engages working memory. Thus, I might have a standard shopping list that I rely on week after week, but I don't constantly have the list in mind. Rather, when I am in the store, I recall the list and use it to guide my search ("Where are the pickles?"). Current models of memory suggest that there is no location where each memory that can be recalled is stored in its detailed entirety. Rather, recalling memories is a constructive process (Schacter 2012). As Schacter et al. note, "Constituent features of a memory representation are distributed widely across different parts of the brain, such that no single location contains a complete record of the trace or engram of a specific experience" (1998, 291–92). If so, then recalling memories to inform action is not the result of selecting among multiple detailed memories, but more likely the selection of a single detailed memory that is constructed from distributed traces of earlier experiences. In many cases, memory-guided action instantiates a one-many, rather than a many-many, behavior space. Memory guided action requires coupling between the constructed memory and possible responses: verbally expressing the memory in recall, using the memory to trigger further thoughts or to guide a specific bodily response, and so on.

The point then is that the same structure is present in both memory- and perception-guided action, a selectivity that is tied to the coupling of an input to a response, amounting to a solution to the Many-Many Problem. Think of visual search as when you visually sift for a visible object in order to respond to it, and think of memory recall as when you mnemonically sift for a remembered object in order to respond to it. Both perceptionand memory-guided action entail solving the Many-Many Problem. A natural thought, then, is that since perceptual selection suffices for perceptual attention (attention in perception), mnemonic selection should suffice for cognitive attention (attention in thought). In James' articulation of the folk-conception of attention, he drew a parallel between the objects of attention in perception and the trains of thought in cognition: "[attention] 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." Recent psychological theories also have emphasized this link between memory and attention (Chun, Golomb, and Turk-Browne 2011; De Brigard 2012).⁷ Given the common structure between perception-guided and memory-guided action, the idea that selection for action is sufficient for attention only when it is perceptual seems an unmotivated restriction. Rather, the common structure of both forms of selection for action suggests attention. To capture this, I drop the restriction to "perceptual" in the expanded sufficient condition to obtain:

 (S_{emp2}) If S selects X for bodily actions, then S attends to X.

Since the relevant Xs are determined by the inputs to a behavior space where those inputs are targets of some appropriate mental state of the subject, then the target of attention will be determined by the intentionality of the subject in action, namely by the intentionality of the input mental state. So, if the subject is visually focused on the ball in a way that sets the parameters for a response such as locating the ball, then the subject is attending to the ball. If the subject is cognitively focused on a memory of a past event so as to inform a judgment, then the subject is attending to that event.

The final step is to ask why one should restrict the condition to bodily actions. Why not include mental actions as well? Mental actions are tied to behavioral outputs that are within the mind, actions like reasoning, imagining, and recalling. Such actions can be prompted by other memories or cognitive states and also by perception. For mental actions, the Many-Many Problem is also present. Here are some examples. I ask you to imagine a "unicorn cow," so you form a distinctive image in your mind that draws on a visual memory of a cow and a visual memory of a unicorn horn, melding the two together. I ask you to solve a difficult logic problem, which requires you to recall the relevant logical inference rules, as opposed to the rules of chess, and to apply them in reasoning out the solution. Or you are presented with two options for dessert at the buffet, both of which look good, but one of which is "healthier". What decision will you make? What considerations are relevant? In these tasks, any mental action exhibits task-relevant selectivity in terms of coupling input materials to a mental response. When those materials, whether from perception or thought, inform bodily action, (S_{emp2}) entails attention. It is hard to see why the same selectivity would fail to entail attention simply because the type of output is mental, rather than bodily. If one's perceptual selection of a ball suffices for attention when one kicks the ball, why not when, on the same basis, one thinks about the ball? To perform mental actions, an agent can select perceptual or mnemonic inputs to guide the output. So, there is selection for mental action as well. This suggests dropping the restriction to bodily actions in S_{emp2} and focusing on actions, both bodily and mental. Accordingly, this yields the most general sufficient condition:

(S) If subject S selects X for some action A, then S attends to X.

Since selection for action is explicated in terms of coupling, where there is coupling to solve the Many-Many Problem, (S) holds that there is attention.

3.5 Personal versus sub-personal selection for action

(S) links selection to attention. Yet just as there is much non-attentional selection in the world, there is also much non-attentional selection for behavior in the world. Consider a Mars rover that has been programmed to respond in certain ways to various inputs. For this simple system, there is a concise representation of its behavioral possibilities as described by its engineers. When the system behaves, it exemplifies selection of input to inform output. This is what makes it natural to speak of the rover as acting. Still, one might be disinclined to ascribe attention to this system, even if it exemplifies a type of selection for behavior. Similar points can be made with simpler machines such as a candy machine that returns treats for change. Thus, for many artificial systems that exhibit behavior, there will be an appropriate behavior space for that system, but typically no attention.

These systems are not covered by (S) since (S) emphasizes that it is the subject that selects the relevant Xs for action. Attention, agency, and mentality are personal level phenomena, in contrast with the sub-personal level. As I use the term, "personal" refers to a certain category of states or capacities that can be attributed to subjects, namely, mental states and capacities. When I refer to the sub-personal, I focus on states or capacities that are not attributed to the subject, but to a part of her, such as her brain. The sub-personal states I focus on are those that have an important connection to personal-level mental states and capacities but do not entail or necessitate the subject's involvement. For example, that the subject's retina registers a single photon does not entail that the subject registers it, say, has a visual experience of a brief blip of light. The retina's registering the photon would be sub-personal, the activity of part of the subject. Under certain conditions, the retina's registering light does causally necessitate that the subject registers it as well, namely, has a visual experience. Here, the sub-personal would be a causal factor in a personal-level state.

The behavior space that is the basis of the Many-Many Problem is a psychological space, where this implies that the subject takes some psychological attitude towards the input. In this way, selection for action begins with a personal-level state, say the subject's perceiving a target or the subject's remembering an important detail. Yet couldn't action be guided only by sub-personal states or states outside of the personal, i.e., states that do not implicate the subject in question? There are cases of a subject's behavior being guided by states that do not implicate that subject and where most would agree that a substantive sense of agency is abolished. Consider cases where you are hypnotized and then made to do things. Here, you are put under the control of another subject's personal-level states, e.g., her intentions, and thus your behavior is guided by states that do not implicate you. While you in some sense act, you are not, in another sense, an agent. You are not in control, but are being controlled. The question is whether control by only sub-personal states of the subject can guide action without abrogating control.

I believe that the emphasis on the personal level is obligatory to keep hold on the idea of action as an agent's doing something. This philosophical conception of action is of the activity of an agent who exemplifies a distinctive kind of control, the sort of control that is abolished in the hypnosis case. Consider, then, a case in which an action is guided by a perceptual state. Now, assume that the perceptual state that guides action is sub-personal, namely a state of the perceptual system that does not entail or necessitate that the subject perceives (recall the retina example). This is behavior that is, as in the hypnosis case, fully guided by states that do not implicate the subject. The subject is absent in the initial perceptual registration of input, and, thus, the action is not guided by the subject but only by some part of the subject's brain. I submit that this situation is not relevantly different from the hypnosis case. In both, a state that does not implicate the subject's involvement is in control of the subject's behavior. While it is true that in the sub-personal case, the state is of the subject's own body or brain, that physical fact does not suffice for restoring the missing element, namely, the subject's own involvement in guiding her action, given how she perceives things to be. Such guidance requires that the input state that guides behavior also reflects the subject's own take on the input. Instead, in the case imagined, the subject moves only as a result of how her brain registers the environment to be while she herself does not register the environment. Qua subject, she is removed from the guiding state, and for that reason, is not exerting control in her action. Yet the starting assumption is that this is a case of perception-guided agency, a way the agent exerts control in action. Similar points can be made for memory-guided actions, as well. Consequently, the agent must be implicated in the input that is coupled to the response for the relevant behavior to count as agency in the thick sense that differentiates agents from candy machines, Mars rovers, and hypnotized individuals. This is why I emphasized that the behavior space must be a psychological space: the inputs imply that the subject bears an appropriate psychological attitude towards them, whether perceiving, remembering, or

thinking. Thus, in coupling an input to an output, a necessary condition on the subject's involvement is secured at the get-go.

3.6 Is attention necessary for action?

In this section, I shall argue for the following conditional, which expresses that attention is necessary for action (ANA):

(ANA) If a subject S performs action A in response to X, then S attends to X for A.

The idea of "response to X" is to be construed as an active response to something that serves as an input to the process of acting. Standard cases are moving in relation to an X (e.g., walking towards it or avoiding it), manipulating an X (e.g., picking it up or throwing it), or having thoughts with regard to X (e.g., wondering if it is edible or committing it to memory). While action with respect to an X might not exhaust the category of actions (perhaps some actions are "targetless"), ANA covers the bulk of mundane actions and, indeed, the sorts of activities of relevance to demonstrative thought and epistemology, activities that can be characterized as directed to an X (e.g., one thinks about a perceived X, one reasons about X, one introspects X; cf. Chapters 7 and 8).⁸ Two assumptions drive the argument for ANA: (1) a pure reflex embodied in a behavior space with a single behavioral path is contrary to action; and (2) the subject's selection to guide action is sufficient for attention. I have already argued for the latter claim in arguing for (S), so let us turn to the former assumption.

Consider what the *absence* of a Many-Many Problem for a subject's behavior would mean. An absence of a Many-Many Problem entails that any behavior generated did not occur in a behavior space requiring selection. This implies that there were no additional behavioral paths beyond the one path taken (this includes the path of not acting). Thus, the behavior space consists of a simple one-one mapping from target to response. All the creature could do was to act on one target in one way. This, however, is just a *reflex*. In particular, to emphasize the very rigid structure at issue, I'll speak of this reflex as a pure reflex. It is not clear that normal human reflexes are pure reflexes in the sense just noted, but it is the latter that are of concern to us.

Reflexes are, of course, forms of behavior and often useful, as when the physician taps your knee. Yet one contrasts reflexes with actions. On being asked why one's leg shot out after the doctor tapped one's knee, one might say, "I didn't do that! The doctor made me do it!" The assumption here is that reflexes contrast with action, though both are forms of behavior (behavior is the genus, reflex and action are the species). If one contrasts pure reflex with action, and pure reflexes eliminate the Many-Many Problem, then to make action available to a subject, the behavior space must be more complicated than a simple one-one map. In other words, for agency to be possible, there must be behavioral options, even if it is just the option of not acting. A behavior space must open up with more than one path, and now there is a Many-Many Problem and the possibility of selective coupling.⁹ All actions then emerge from an appropriate behavior space where actions entail selective coupling, choosing one among other behavioral paths. This "choice" is not available in pure reflex.

This suggests a certain picture of agency that emanates from the Many-Many Problem: agency implies a solution to the Many-Many Problem, a solution that entails the coupling of input to guide a behavioral output in a behavior space that presents the agent with options. Since this traversing of a path involves a form of selection for action in input-output coupling, and such selection is sufficient for attention given (S), then action implies attention. Specifically, action on some X, whether a thought one entertains, a memory one recalls, or an object one manipulates or otherwise acts on, entails attentional selection of X. This claim will suffice for certain philosophical applications of attention in later chapters. Perhaps surprisingly, attention is not something that merely aids action in respect of an X or makes it more efficient. It is not something that is optional for acting in this way. Rather, attention, without coupling, no such action would be possible.

3.7 Is selection for action necessary for attention?

Let us now consider (N):

(N) If S attends to X, then S selects X for performing an action A.

Why think that (N) is true? Let's soften the ground a bit first and note that in standard cases, wherever one finds attention, one also seems to find selection for action. Thus, I attend to a conversation so as to verbally shadow it, to listen to it, or to prepare a response; I attend to an object to catch it, to follow it, or to get away from it; I shift attention around to locate an object, to find a hiding place, to locate the shortest path to my destination; I attend to a line of thought to figure out what is right, to locate a solution, to prepare an answer; I attend to a memory to recall what is important, to enjoy a fantastical image, to guide my shopping. Here, selective attention to some target is for the purpose of performing a task, and this performance implies selection for action. This suggests that (N) is plausible.

I do not see any way to derive selection for action from the concept of attention as the latter is the target of analysis. Rather, my strategy will be an inference to the best explanation. Psychological states, like perception, memory, and thought about an item, do not on their own conceptually entail attention to that item. In later chapters, I will consider views where conscious perception implicates attention as part of a necessary causal process, but unconscious states do not in any way implicate attention. Suppressed unconscious desires or memories do not entail attention, though they can engage attention, perhaps with the help of a therapist. Similarly, unconscious perceptual states of an object, such as those in blindsight, do not entail attention to the object. Something further is needed if attention is to emerge from a subject's mental states. Given James' explication of attention, a natural response is that some form of appropriate selection is required over and above being in a mental state. There are two main options: selection for action and selection for consciousness. In the remainder of this chapter, I defend the selection for action account by responding to the two main objections to (N). This is to establish that selection for action is a plausible proposal of what one must add to a mind to yield attention: without it, there is no shift from mental states like thought or perception to attention. In subsequent chapters, I shall argue for the selection for action account as the best explanation by showing that selection for consciousness does not track attention in the required ways (see Section 4.2 on unconscious attention).

3.7.1 Attentional capture

An obvious challenge to (N) is stimulus-driven, bottom-up attention. Can the selection for action account, which endorses (N), account for these cases? Attentional capture and the phenomenon of pop-out present potential cases. Chapter 1 revealed, however, that many cases of pop-out are in fact goal-directed, so one's attending to a specific target defined in task instructions is subject to control as defined in Section 1.7. In that way, pop-out is not the challenging case. Rather, it is attentional capture, fully automatic attention that can disrupt an agent's current goal-directed behavior (recall the two neural networks for attention discussed in Chapter 1). The question, then, is whether a selection for action account of attention can allow for fully automatic attention, and thus whether:

If S's attention is captured by X, then S selects X for performing action A.

The answer is yes, but one must make explicit the relevant conditional:

If S's attention is captured by X, then S automatically selects X for performing action A.

Automatic selection of X for action is, of course, a type of selection for action. Given the definition of automaticity in Chapter 1, to say that the selection of X is automatic is to say that selection of that target is not intended. Yet that is precisely what happens in attentional capture, at least initially: the capture of attention involves the capture of action.

Here is the basic structure of my argument with respect to attentional capture. Attentional capture is driven by cases where there are sudden changes in one's mental state in respect of an object (or feature): a loud sound (or a sudden shift in the pitch of a current sound), a swooping bird, a fragrant smell, a pleasurable memory, a twinge in one's calf, or a disturbing thought. One thereby moves from not having any mental states directed at the object (which has not yet appeared) to having a mental state directed at that object. Yet, since none of the imputed mental states directed at the object on its own entails attention, something else must be added for attention to appear. For example, in the perceptual cases, the appearance of a sudden stimulus yields only the perceptual registering of a change, a transition from not perceiving an object to perceiving it. But the perceptual registering of a change does not entail that one attends to the object at issue. In general, there can be alterations in a subject's mental states without these alterations engaging attentional capture.¹⁰ Something else is needed to allow for attention, but what? Proposal: selection for action. The input state must be coupled to action.

This proposal that selection for action transforms the mere registering of change to attention fits with a common comment in psychology about attentional capture, namely, that it is tied to an orienting response. This suggests an action that is automatically driven by the suddenly appearing stimulus, an input-output coupling that can be contrary to the subject's current intention. Where there is the capture of attention, what captures attention always engages a response: you move your body towards the object, orienting towards it; you think about it; you explore it with the other senses. In such cases, the resulting actions are automatic, at least initially, because none of them are intended. Still, to the extent that this automatic response yields information that is goal-relevant, control can kick in. One subsequently intends to explore the object further. Attentional capture gives way to controlled attention.

Similarly, thoughts and images can also pop into your head as an object pops into your sensory field. The same considerations apply here as well, though in the mental case there can be a small distance between input and output, e.g., while the input might be the flashing of a specific visual image, the output would be the maintaining of that very image. Input and output are nearly identical, the latter simply involving a response to the image. The first point is that the flash of a thought is no different than the flash of a change in perception. Both yield alterations in what is given in one's mental states that do not entail or necessitate attention. It is only when thought and perception engage with the mind in a further way that one moves towards attention. Thus, consider the flash of a thought that leads one to ponder it further. An embarrassing image of last evening's faux pas might involuntarily flash in one's head, but, if the thought amounted to nothing more than that, no attention emerges. There is just the fleeting image. It is only when that image engages further activity, when one ponders it, laments it, or just sustains the image as one internally cringes, that attention emerges. The point then is that, in attentional capture, attention enters the scene when the item that does the capturing not only alters the shape of one's mental states, but does so in a way that engages a response. Otherwise, there is only the mental registering of a change but no attention to it. The point, then, is that attentional capture is more than one's mental states changing in response to a sudden stimulus. Attention is not just a shift in consciousness or an alteration in one's mental states given a new stimulus. Rather, it is only when this change in one's mind engages with something further that attention comes on the scene. Without this further ingredient, attention is not present. The proposal is that this ingredient is selection for action: the change engages a response.

3.7.2 Attention as vigilance

In the Introduction, I noted the possibility of a pluralist account of attention, that attention is many things. Indeed, many theorists of attention identify both vigilance as well as arousal as distinct forms of attention that do not involve selection (e.g. Robbins 1998). The question, then, is what these notions come to. In a recent review on vigilance, Oken et al. (2006) note at least three meanings of "vigilance" in the empirical literature: (a) an ability to sustain attention over a period of time; (b) in animal behavior, attention to potential threats or dangers, including hyper-vigilance; and (c) a narrower notion that refers "to arousal level on the sleep-wake spectrum without any mention of cognition or behavioral responsiveness" (p. 1885). Indeed, echoing concerns about attention noted in the Introduction, the authors conclude that "the field has been hindered by inconsistent or poorly defined terminology ... [A]voidance of the term vigilance because of its varied definitions would be most helpful" (p. 1895). Similar points can be made about "arousal".

These observations present a challenge: if there is a distinct conception of attention that is tied to vigilance, then there is further work for conceptual analysis. It is worth noting, however, that in Oken et al.'s characterization of two of the three notions of vigilance, the notion attention is used. Vigilance seems to be a way of attending, and then one might ask whether the conception of attention used to explain vigilance is just selection for action. If so, then vigilance will not be a separate form of attention, though it might be a way or manner of attending. In a review of the literature, Parasuraman et al. (1998) focus on sustained attention in order to detect infrequent events. They cite pioneering experiments by Norman Mackworth where subjects were involved in a two-hour task where they had to monitor the movement of a hand on a clock, identifying double jumps (3-5% of the time) against a background of frequent single jumps. Mackworth, and others, observed a characteristic vigilance decrement, both in terms of a decrease in detection rate and an increase in reaction time over the course of the experiment. Vigilance fails because the subject is unable to sustain attention.

If a change in vigilance is typically measured by the vigilance decrement, and the latter is tied to certain properties of task performance, namely changes in detection rate and reaction time, then it looks like vigilance, even if it is different from selection from action, supervenes on it. That is, changes in vigilance are measured by changes in selection for action as measured by behavioral outcomes. This suggests that vigilance is a property of selection for action over time: vigilance is a measure of how effective subjects are in selecting for action. One assesses vigilance by assessing how subjects sustain selection for action in an experimental setting.

What then of hyper-vigilance, which an animal might deploy as it stays attentive for predators? One reason to think that this sort of vigilance does not require selection is that vigilance involves keeping an eye out for, but not selecting, any predators (they aren't there yet). When a predator appears, fleeing, not hyper-vigilance, is needed. Thus, there is no object that is being selected for action at the time of vigilance. Still, selection for action underlies hyper-vigilance. Consider two chutes from which a ball can be launched, and the difference between readying yourself to catch the ball from one chute and readying yourself to catch the ball from the other. These are distinct states of readying yourself for action, guided by the information you receive from the environment, such as the location of the chutes. Vigilance involves keeping the eyes fixed on the relevant chute, and this is a motor action, one where the agent perceptually attends to a location or object (the chute) and uses this to sustain a motor response, in this case to keep the eyes locked on the target. One can increase vigilance over the two chutes by rapidly shifting the eyes from one to the other, maintaining a state of readiness to catch the ball. This is plausibly hyper-vigilance. Similarly, consider a hyper-vigilant gazelle drinking at a water hole, a dangerous thing to do since the eyes are not scanning the environment for danger. Hyper-vigilance is exemplified by the gazelle's constantly raising its head to scan the environment, disrupting continuous drinking. Vigilance is often invoked as a form of attention, but on the best current understanding of vigilance, it looks to be a higher order property of selection for action, a manner of this form of selection.

Let us return to (N):

(N) If S attends to X, then S selects X for performing an action A.

I have noted anecdotally that when focusing on cases of attention, one finds selection for action. Further, I have considered two counterexamples to (N) where attention might not seem to entail selection for action, but I have suggested that selection for action is plausibly present. To fully secure (N) one will have to consider whether attention entails consciousness, but let us, for the moment, take (N) as plausible and flesh out the selection for action account.

3.8 Attention as selection for action

I have built an initial case for (S) and (N). Sympathetic readers might then consider the biconditional expressing selection for action (SfA):
(SfA) S attends to X if and only if S selects X for action.

The proposal is that this action-centered approach provides the best current answer to the metaphysical question: What is attention? I began by extracting the empirical sufficient condition from experimental practice in the science of attention and then expanding this condition to (S). I then argued for (N) by defending it from the most obvious counterexamples. Accordingly, the answer to what attention is focuses on attention's ties to action. One could make the bold move from SfA to an identity claim (I):

(I) S's attention to X is S's selection of X for action.

One could treat this as a hypothesis whose value will be revealed by whether it can be theoretically useful. Certainly, given that (I) and (SfA) imply (S), and given (S)'s foundational role, as discussed in Chapters 1 and 2, via the empirical sufficient condition, the action-centered approach does facilitate understanding of attention. Let us, however, take the status of (I) and (SfA) as the best current hypotheses. The question now concerns the specific shape of the answer to the metaphysical question, given (I).

The specific details of the selection for action account depend on how one interprets the phrase "selection for action." Specifically, how should one understand "selection"? Using a behavior space, one can pictorially represent two options: (a) selection is the coupling, the arrow that ties input to output; or (b) selection is the specific input that gets tied to the output:



Figure 3.2 Should attention as selection for action be identified (A) with the perceptual state of visually experiencing the ball, or (B) with the process of coupling that experience to the kicking movement?

Put in words, selection for action is either the complex process of tying an input to the output, or it is just in some sense the input itself when that input serves an output. I suspect that many theorists of attention sympathetic to the selection for action approach will opt for a process-based answer to the metaphysics of attention and thus opt for tying selection to coupling. The other option might seem somewhat strange, but let me say a few words in favor of it.

Attention is a personal-level state that exhibits intentionality. One attends to an X. The intentionality of attention must be explained. A parsimonious selection for action account will explain the intentionality of attention by appeal to the intentionality inherent in a behavior space when the Many-Many Problem is solved. What one is looking for is a mental state that has the same intentional directedness attributed to attention. Let us imagine that there is a single ball to be kicked. Where one attends to the ball to kick it, what relevant states in behavior space are similarly directed? The answer is clear: there is a visual state that is also directed at the ball. The subject has a visual experience of the ball. The initial thought, then, is that attention to the ball just is the experience of the ball when that experience meets an additional condition: it is coupled to an action. Where a subject is visually experiencing a ball and that experience guides an action by contributing to parameter setting for the action, the subject is attending to the ball. Failed actions still allow for attention so long as coupling is initiated. In general, facts about coupling determine the subject's state of attention.

To flesh out the proposal a bit, consider a case where a subject visually experiences multiple targets. So, if one sees a basketball and soccer ball, the visual experience represents both objects. The relevant input in respect of coupling when one kicks the soccer ball is, at least in part, the specific visual content concerning the soccer ball. That content identifies part of the intentionality of the experience, and it is that content which also contributes to parameter setting of the action of kicking the soccer ball, say by identifying the target of action. When the agent kicks the soccer ball, the relevant input-output coupling is explained in terms of one's experience of the soccer ball and not one's concurrent experience of the basketball. It is the former content that sets the parameters for the kick. The reverse would be true were the agent to kick the basketball instead. Thus, a specific aspect of perception explains coupling, namely, the relevant perceptual content that contributes to parameter specification. Accordingly, the intentionality of the agent's attending to the soccer ball is tied to that aspect of perceptual content that accounts for the experience's role in coupling, a content that

also accounts for the intentionality of the experience, i.e., what it is directed at. So, the kicker attends to the soccer ball because his experience of the ball is coupled to his response. In general, the relevant aspect of the personal-level input state that sets parameters for coupling fixes the target of attention, i.e., what the subject is attending to. The subject attends to a perceived feature when that aspect of perception is coupled to his response; he attends to a thought because his thinking that thought is coupled to further trains of reasoning; he attends to a memory because his remembering an episode is coupled to further responses to that past event, and so forth.

While the analysis of attention appeals to coupling, attention on this version of the selection for action account is not identical to coupling.¹¹ Think of a simple causal functionalist account of a mental state that identifies what that state is in terms of its causal role, its input-output profile in terms of causes and effects. Even though the analysis of what that state is mentions those causes and effects, a token or instance of that state is distinct from the specific causes and effects associated with it. Thus, if pain were identified with that state caused by As and which itself causes Bs (whatever As and Bs are), a state of pain is nevertheless distinct from its causes and effects. Similarly with attention understood as selection for action. A subject's attention or attending to X is to be understood as a selective state with respect to X that is individuated as attentional selectivity because the selectivity is for action. What is attention? It is just a psychological state directed towards an input when that state's being so directed, given its content, is coupled to action in solving the Many-Many Problem. In that sense, attention is selection for action.

The contrast to this input-centered approach is a process-centered approach that identifies attention with the coupling that links the input state to an output. Those who identify attention with this process need not deny that attention has intentionality. They can take on board the account of the intentionality of attention just noted, but they will add that the relevant input state is part of a complex process of bringing a target to bear on action in response to the relevant content of that input state. It is the process of tying the input state to action that identifies attention. Accordingly, attention is a process of selecting an input to guide an output, where it is the input that identifies the intentionality of the subject's attentional state. There is no need to resolve which of these two options provides the best elaboration of selection for action, and it might be that both proposals will have their place depending on whether one focuses on the subject's state of attention or on the subject's attending. The former emphasizes a state of mind directed at the world while the latter suggests a type of activity.¹² Accordingly, the input-centered selection for action account might work better in an account of attention as a state of mind while the process-centered account might work better in an account of attention as an activity. That there are options here reflects the different meanings and uses of "attention" and "selection" (cf. attending, selecting, selectivity, attentiveness). What I hope to have shown, however, is that the selection for action account is neither idiosyncratic nor established by fiat. It has its roots in shared assumptions in experimental practice and provides a promising way to answer the metaphysical question.

3.9 Attention as cognitive unison

Christopher Mole (2011) has recently presented a detailed and distinctive theory of attention. His account emphasizes attention as attentiveness understood as a feature of action, namely the manner in which the agent's activity unfolds. Accordingly, he presents his theory as an adverbial account of attention. In this book, the central question is the metaphysical question:

1. What is attention?

In contrast, Mole prioritizes the question:

2. What is it for something to be done attentively?

In part, Mole emphasizes (2) over (1) given his stance against theories of attention that identify it as a process (see Chapter 2.7 on Mole's argument). I shall set the issue of priority aside in favor of noting that a complete theory of attention answers both questions. If Mole's theory answers (2), one can conjoin it with the selection for action account's answer to (1). Indeed, both theories share an action-oriented approach to understanding attention. Given that Mole's theory can contribute to a complete account of attention, let us briefly consider it.

Whenever one performs a task, one brings to bear various capacities that are deployed in unison to serve the task, like the members of a wellrehearsed orchestra playing a symphony. Acting attentively is just using one's resources in this unified manner. This leads to a simple account of acting attentively:

Unison: A subject *S* performs task *T* attentively if and only if *S*'s performance of *T* displays unison.

The initial problem with this proposal is that all tasks come out as performed attentively: resources used to guide any task are brought together to serve that task (see Mole, p. 64). Yet, some actions are done inattentively.

Mole's theory differs from the simple account of unison in two ways. First, he restricts the relevant resources to specific cognitive resources. Second, Mole's account is presented in *negative* terms. It defines what must not happen if there is to be attentive performance. Mole then provides the following definition of attentive action in terms of cognitive unison:

Let α be an agent, let τ be some act that the agent is performing, and call the set of cognitive resources that α can, with understanding, bring to bear in the service of τ , τ 's "background set".

 α 's performance of τ displays cognitive unison if and only if the resources in τ 's background set are not occupied with activity that does not serve τ [I have replaced "task" with "act"].

This is complicated, but the central idea is that for an action to be performed attentively, every relevant cognitive capacity is put to work to aid the action or, if not aiding the action, it is at least not put to work in the aid of any other action. One can appreciate the essential feature of the account without explicating the central ideas, but a few words about cognitive processes, the background set, and the notion of understanding. The theory explicates cognitive processes in connection to the agent's propositional attitudes such as beliefs, desires, and intentions. Specifically, cognitive processes operate over representations that can determine the content of the subject's propositional attitudes. Thus, in doing sums, one calls on cognitive processes of relevance to arithmetic, and these processes put one in a position to think about, ponder, or recall numerical operations. The set of cognitive processes that can serve an action form the background set for that action. These are capacities that the subject can deploy with understanding in that the agent can articulate what she is doing, how she is doing it, and for what reason. Such understanding plays a specific causal role: it controls and guides behavior. Thus, in completing an academic paper, the agent can describe, explain, and justify what she is doing, and she thus expresses a state of understanding that guides her action. In sum, an action is done attentively when all the relevant cognitive capacities that constitute the action's background set are deployed by the agent with understanding to serve the action or, if those capacities are not deployed, they are not deployed to serve any other action.

For Mole, unison entails that "the resources in τ 's background set are not occupied with activity that does not serve τ " (51). This requires "the *absence* of any irrelevant processing" (70). There are, however, two ways that the absence of irrelevant processing can be achieved. First, *all* the members of the background set are in sync, working towards the same task. Second, some of the members are working towards a task while the others are dormant. In both cases, irrelevant processing is absent. Mole, however, notes that "the cognitive resources of a normal waking brain tend to find something to occupy them, and so one can typically expect there to be some processing going on in any given resource" (70). Where processes in the background set are dormant, there is the danger of their being distracted, siphoned off to serve irrelevant processing.

It is useful at this point to contrast a positive version of cognitive unison that is not bothered by distraction. On this account, one might hold that it is enough for cognitive unison that some of the relevant capacities are deployed in unison for the task. Thus:

 α 's performance of τ displays cognitive unison if and only if *at least some* of the resources in τ 's background set are occupied with activity that serves τ .

Distraction, i.e., the siphoning off of some of the members of the background set, is not a threat to cognitive unison, so long as other members of the set are operating in unison to serve a task. In general, attentive performance seems less fragile on the positive account than on Mole's negative account, for on the latter, if just a single member of the background set gets distracted on matters that are task irrelevant, performance is no longer attentive. No such consequence follows from the positive theory.

In fact, Mole's theory can be put in positive terms: to say that no member of the background set is serving tasks other than the targeted task is to say that if any member of the background set is serving a task, then it is serving the targeted task. Rather, the difference between the two versions is best stated in terms of quantification (Mole speaks of this as the difference between absolute and threshold accounts of unison). The fragility just noted is really a result of whether a cognitive unison theory universally or existentially quantifies over the background set. That is, whether the theory either requires some or all of the members of the set to be operating in unison. On the universally quantified theory, unison is fragile: one bad apple spoils the lot. Existentially quantified accounts can allow for bad apples, to a point. Mole's main objection to the existential account is that it seems that (nearly) all tasks will be attentive (personal communication). For any task will involve unison in the existentially quantified sense: there will be some number of capacities deployed for that task. Since one must allow for inattentive tasks, this seems to be a *reductio* of the existential account. In response, one might opt to speak of degrees of attentiveness and inattentiveness. One is more or less attentive, and to the extent that one can individuate the capacities that belong in a background set, then one can speak of more attentiveness when more of the background set is involved in the action; less attentiveness when less is involved. The same goes for talk of inattentiveness. This leads to an objection to Mole's universally quantified version of the theory, for, strictly speaking, it does not allow for partial or divided attention. Mole notes several ways where his version of cognitive unison can allow for talk of partial attention, but, in the end, he emphasizes that just as something that is only partially built fails to be built, something that is only partially unified fails to be unified.

If one thinks that there is partial attention, then this might give one reason to prefer the existential formulation of cognitive unison. But it is worth raising a different question: is there really an interesting notion of partial attention, of attention itself being more or less?¹³ That is, can attention really be scaled in this way? It is certainly true that, pretheoretically, it is natural to speak of degrees of attention, but is this common talk something a theory of attention should take on board, rather than jettison? Should one think of attention like energy or a resource that subjects can tap into more or less? Can such metaphors be given a firm basis?

In the climactic gun fight in Sergio Leone's The Good, The Bad and The Ugly, the three gunslingers stand equidistant from each other along the circumference of a large circle, poised to shoot. The camera cuts between each individual's face and, as the scene drags on, focuses directly on each individual's eyes that rapidly flick from one opponent to the other. What would be the best strategy for each gunfighter? They might divide attention to each opponent rather than shifting attention from one to the other. Perhaps this is accomplished by splitting their attentional spotlight. But it is likely that in such a life-or-death situation, one would feel more confident in shifting attention back and forth quickly, rather than dividing attention. Why? Because in dividing, one feels there is a cost. This is why it is easy to speak of giving undivided attention. In dividing, one is less attentive, and this can be seen in slower reaction times, which in the case of your average duel at high noon is literally life threatening. It is reaction time decrements that provide evidence for less attention.

The challenge is to make good on the idea that attention can be more or less, something one can quantify. It is one thing to point to the behavioral upshot as more or less (reaction time being something one can quantify), another thing to point to the underlying capacity, attention, as something that also can be more or less. Note that many theorists of attention tie attention to a process, but in what sense is a process something that there can be more or less of? For example, is there more or less combustion, as opposed to more or fewer things that combust, or more or less energy that is produced? Thus, is there more or less attention, as opposed to more or fewer things that one attends to, or more or fewer ways that one can respond to what is attended to? The idea of partial attention poses a potential challenge to unison accounts. But turn the table: what does it mean for attention to be partial? Is there any adequate analysis of this idea that can't be understood in terms of the number of targets of attention or behavioral measures like increases or decreases in reaction time?¹⁴ A defense of Mole's universally quantified theory of cognitive unison is just to deny that attention comes in degrees.¹⁵

3.10 Summary

The link between attention and action has pervaded our discussion of attention. From the beginning of experimental work on attention, theorists have made that connection as expressed in the empirical sufficient condition (Chapter 1). That condition also serves as an organizing principle to isolate attentional effects in the brain (Chapter 2). In this chapter, the link to action has been presented as central to a complete understanding of attention and attentiveness. A promising answer to the metaphysical question appeals to selection for action. The suggestion, then, is that an action-centered account provides the best way to generate a comprehensive theory of attention. Of course, more work needs to be done to assess this approach, but the current claim is that as things stand now, with growing skepticism concerning attention rising in various quarters, the action-centered approach provides the best hope of explaining attention. It should be the working empirical and philosophical hypothesis about the nature of attention.

Suggested reading

Allport (1987) and Neumann (1987) provide foundational discussions of attention as selection for action and are unduly neglected. They deserve to

be read carefully. Wu (2011b) provides a recent defense and elaboration of their ideas, and Wu (2011a) develops these ideas as part of a philosophical theory of agency. The cognitive unison theory is presented in Mole (2011). A different adverbial theory drawing on the semantics of questions is given by Koralus (forthcoming). For an earlier discussion that ties attention to an activity of questioning the environment, see Eilan (1998). Peacocke (1998) has an important discussion of acting attentively that is not discussed in this chapter.

Notes

- ¹ There is experimental work that suggests that even in conditions of presumably low load, say two stimuli, attention to one of the two leads to withdrawal of resources from the other (e.g. Pestilli and Carrasco 2005).
- 2 Carolyn Dicey Jennings has taken a distinctive position on attention, as well as raising some problems for the selection for action account. She argues (2012) that my (2011b) version of the selection for action account conflates the subject with the agent. In this chapter, the notion of a subject is taken in the minimal sense of a bearer of psychological states. On this view, being a subject does not entail being an agent, but being an agent does entail being a subject.
- 3 Generally, we will have many Many-Many problems. There are, after all, different kinds of attentional targets, so, even if a creature is confronted with a single object, that object exemplifies many perceptible properties. So long as those properties do not drive a pure reflex, then the subject must select relevant properties among irrelevant properties at the level of the input. For example, spatial properties may be more useful than color properties in one context and vice versa in another. To react accordingly, one must select one among many features.
- 4 In some recent discussions of the role of vision in action, it has been argued that conscious vision does not directly set parameters for precise motor movements, but rather identifies the target for motor computation, e.g., which object to reach for (see Milner and Goodale 1995; Campbell 2003; Clark 2001). On this picture, conscious vision still sets a parameter for action (it identifies the target) but it doesn't set all the parameters, such as the precise metrical information about the object's structure: its size, shape, and location. Nevertheless, even on this model, the input-output coupling proceeds from conscious vision to movement, even if there are unconscious and subpersonal visual states along the way (for an

argument for there being unconscious states mediating mundane actions, see Wu 2013). The input remains the subject's visual experience of a specific object to act on.

Whether this interpretation of the role of visual experience in guiding action is correct is a complex matter. The issue has generated a large literature in philosophy and cognitive science and concerns the function of two anatomically separate streams in primate cortical vision (what are called the dorsal and ventral streams, the former projecting into the parietal lobe, the latter into the temporal lobe). The original hypothesis about the two primate visual streams as what and where streams, respectively, was given by Leslie Ungerleider and Mortimer Mishkin (1982). They claimed that the dorsal stream serves spatial awareness (where) while the ventral stream serves visual object recognition (what). David Milner and Melvyn Goodale have argued for a different account of the streams as serving perception and action, a case they systematically argue for most recently in their (Milner and Goodale 2006). For an engaging overview of the issues, see their (Goodale and Milner 2004). For some recent empirical queries about Milner and Goodale's influential account, see (Schenk and McIntosh 2010). For relevant philosophical discussions, see (Mole 2009), (Briscoe 2009), (Brogaard 2012) and (Wu forthcoming) among others.

- 5 Thus, there was no controversy in the history of modern psychology when investigators switched from dichotic listening, with its selection of specific verbal streams for verbal shadowing, to visual search, with its selection of visual objects to identify a target. This suggests a broader notion of task in the sufficient condition than something restricted to specific experimental paradigms. What makes the incorporation of new experimental tasks plausible in investigating attention is precisely that the tasks identify targets as task-relevant.
- 6 Alan Baddeley firmly established modern research on working memory with his influential model (Baddeley and Hitch 1974) that contained three components: (1) a central executive that controls selective attention; (2) a visual-spatial sketchpad; and (3) the phonological loop. The phonological loop is itself composed of two sub-parts, a *phonological store* that keeps representations of sounds, and the *articulatory loop*, where items in the phonological store are actively rehearsed, e.g., in subvocalization (for example, many of us can remember an item by repeated internal articulation: "buy juice, buy juice, buy juice"). Baddeley also focused on visual information in the visual-spatial sketchpad although he allowed a role for the

sketchpad in integrating information from other sensory modalities and from long-term memory. Finally, the central executive was, by Baddeley's own acknowledgement, something of a *homunculus*, a black box into which theorists sweep items yet to be explained. In recent work, Baddeley has added a fourth component to the model: the *episodic buffer* (for a review, see Baddeley 2012). This component now takes the role of integrating information from sensory modalities and long-term memory.

- 7 The connection between long-term memory, working memory, and attention has been gaining much traction in recent years. Indeed, one conception of working memory just is a conception of long-term memory that is currently in the focus of attention (Cowan 1995).
- 8 I consider actions that are responses to and directed at targets, but what of actions without targets? The idea is that sometimes one doesn't act in respect of an *X*; one just acts. In light of these cases, I am tempted by the claim that attention is necessary for "targetless actions" (ANTA):

(ANTA) If S performs action A, then S attends to some X for A.

The reason why I am inclined to endorse ANTA is that the actions at issue really fall under ANA: actions that seem to be targetless have targets in that they are responses to various inputs. These are in fact complicated issues that require more discussion than can be given here (ANA will suffice for the purposes of later chapters). Let me just give one example that is often brought up in this context to give a sense of the issues: meditation. Part of the challenge of developing this line of thought is to be very clear about what sort of action is involved here. After all, at some point, meditating might lapse into not doing anything (one falls asleep), and then the issue does not engage. But in cases where getting into a meditative state is active, something one is doing, then there is a relevant behavior space: one begins with an input state that determines the path needed to arrive at a meditative state. If my mind is "full of thoughts," then that initial state programs a response that requires the emptying of those thoughts, with the ideal output being a state of mind that is "thoughtless". To the extent that maintaining this meditative state is active, then there is a path that takes that state as input in order to program a response that is the maintaining of that state. Notice that many aspects of the action can become *automatic*, as in the case of advanced meditators. But the automaticity of the path does not eliminate its structure, namely, an input informs an output. Coupling is still present, but it can occur without the subject actively intending to do so. Some forms of meditation involve attention, e.g., in what is called mindfulness meditation. It is sometimes said to me that some forms of meditation are precisely to remove one's attention. I am not sure what that means, unless one has drifted off to sleep!

- 9 A slight complication: The Many-Many Problem doesn't follow just because there is more than a one-one map. You could have two pure reflexes after all. The geometrical structure of the behavior space must also be one that allows for behavioral options.
- 10 I am relying on the claim that one can be conscious of X without attending to X. As we shall see in Chapters 5 and 6, this claim is denied by those who endorse a gatekeeping view of attention: one is conscious of X only if one attends to X. But if the gatekeeping view allows for attentional capture, it holds that the capture of attention brings things to consciousness. On some renditions of this view, attention enables unconscious perceptual states to become conscious once attention is pulled to them (or again, think of therapy that brings unconscious desires and memories to the light of consciousness). But then, this allows for a version of the argument given in the text, except directed at unconscious perception of change.
- 11 We can have weaker but related views, such as that attention is a higher-order state that supervenes on the input state that meets further conditions, namely is coupled to action. But the identity theory is easier to grasp, so I have presented the view in terms of (I).
- 12 It is possible to see the parsimonious approach as a form of *adverbialism* (see next section). The idea is that attention is just a way of perceiving or thinking, namely, when perception and thought are coupled to a response.
- 13 As we shall see, in the empirical study of attention and consciousness in Chapters 5 and 6, it is often claimed that we can be conscious of items outside of attention. Sometimes, the claim is that there is consciousness outside the "near absence" of attention, implying, presumably, that there remains at least the possibility of residual attention.
- 14 A natural response is to appeal to the neural basis of attention as a means of explicating degrees of attention, since neural activity can exhibit gradations. The open question raised in Chapter 2 concerns the relation between such activity and attention. One should not infer from the fact that the neural basis of attention exhibits gradations, say in neural activity, that attention itself exhibits gradations.
- 15 In recent work, Philipp Koralus (forthcoming) has proposed an *erotetic theory* of attention that offers an alternative adverbial account to Mole's

cognitive unison theory (see also Cumming, "The Attentional Foundations of Coherence"). The central idea is that we can think of perception as a way of interrogating the environment, and attention as connected to answering questions. Thus, when I look for my keys, perception aims to answer the question, Where are my keys? Attention is a way of focusing one's attempts to answer that question. This suggests an interesting way to conceive of the function of attention in performing a task, namely to aid in the answering of task-relevant questions. To flesh out this idea, Koralus uses the semantics of interrogatives.

In modeling the relationship between questions and answers in natural language, it has been recognized that in order for something to count as a fully *congruent* answer to a question, what the putative answer is "about" has to match what the question is "about". In English, stress patterns are involved in marking this kind of aboutness, while other languages like Japanese also include special grammatical particles to mark the relevant distinctions. For example, reading the italicized word with appropriate stress, "*Adam* philosophizes" is an appropriate or congruent answer to "Who philosophizes?" while "Adam *philosophizes*" is not a congruent answer to that question, although it is a congruent answer to a different question: "What does Adam do to make the world a better place?" The erotetic theory then emphasizes the following two features: (a) the set of possible answers that individuates a question; and (b) the contribution of *focus* marking in determining congruent answers, analogous to the use of italicization in the previous example of "Adam philosophizes".

Like Mole, Koralus also emphasizes tasks. Let us take as the basic case for the erotetic theory situations where agents monitor the performance of their task so as to query whether the task is completed. This task monitoring can be modeled by a set of questions that the agent poses. Thus, in visual search for a red square, monitoring task completion will involve the question, *Where is the red square*? This question determines a set of answers that specify the *completion conditions* for the task, namely propositions identifying the possible location of the square: {red square is at p_1 , red square is at $p_2 \dots$ }. Now, most tasks are complex, involving subtasks, each with its own completion conditions. For complex tasks, agents will need to monitor subtasks as well, each tied to its own question and set of possible answers that fix its completion conditions. To take another example, in reaching for a mug, certain questions seem appropriate to assessing completion: Where is the mug? Is the arm on the right trajectory towards the mug? Is the grip being prepared of the right sort for a mug? Such monitoring need not be explicit or conscious, but performing a task requires monitoring of this sort.

The central idea of the erotetic theory, then, is that to do a task attentively is to be sensitive to the answers to the questions set by a task. Thus, to visually search for one's keys attentively is to be sensitive to target location, one of the answers to the question set by the task: Where are my keys? The problem, however, as we noted in the semantics of questions, is that not all answers are congruent. To aid congruency, attentional focus plays the same role in attentive action as focus does in interrogatives. Both help to determine congruent answers. Doing a task attentively, then, involves the agent's being sensitive to congruent answers to the questions posed by the tasks. More simply, attention helps to get us to the right answers.

4

ATTENTION AND PHENOMENOLOGY

4.1 Introduction

Consciousness does not play any role in the selection for action account of what attention is, yet William James emphasized that "focalization, concentration, of consciousness are of [attention's] essence" [my italics], and many will agree, thinking of attention as necessarily conscious. In this and the next two chapters, I shall consider answers to the metaphysical question that take attention as essentially tied to consciousness. This chapter examines attention as a distinctive form of consciousness whereas Chapters 5 and 6 focus on attention. Such consciousness, or what I call the gate-keeping view of attention. Such consciousness-centered approaches to the metaphysical question for action account.

Indeed, the common visual metaphor of the attentional spotlight echoes James's talk of focalization and concentration and suggests a different answer to the metaphysical question: attention is a distinctive way of being conscious. This yields a phenomenal conception of attention: there is something it is like to attend, a distinctive phenomenology of attention.¹ Accordingly, what it is like to visually experience an object when attending to it is different from what it is like to visually experience the object when not attending to it. The phenomenal difference between these experiences points to the distinctive

phenomenal quality of attention. Some speak of spotlights and zoom lenses (Eriksen and St James 1986); others speak of highlighting (Campbell 2002). This phenomenal conception, couched in talk of conscious attention, is called upon to do philosophical work in discussions of demonstrative thought (Chapter 7) and justification in epistemology (Chapter 8). A central question to be addressed in much of this chapter is this: is there a viable phenomenal conception of attention—a what-it-is-like to attend—such that attention is revealed as a distinctive mode of consciousness?

To focus ideas, consider two questions:

- 1. Is phenomenal consciousness of attention's essence?
- 2. Is there a characteristic, uniform way it is like to attend to something?

Now, if there is unconscious attention, then the answer to (1) is "no". Even so, (2) remains a substantive question, albeit restricted to those cases where attention does have an associated phenomenology. There are, then, four possible positions depending on how one answers (1) and (2). A natural version of the phenomenal conception answers both questions affirmatively: attention is always conscious and has a uniform phenomenology as a distinctive mode of consciousness. All plausible phenomenal conceptions, however, answer (2) affirmatively, so the idea of a uniform phenomenology of attention will be the main target of critical reflection. The difficulty, I hope to show, is that there is a dearth of detailed proposals as to what the phenomenology of attention might be.

Section 4.2 addresses question (1) by presenting empirical evidence for unconscious attention. Section 4.3 examines recent empirical work on the phenomenology of visual attention. This work in fact does not reveal a distinctive phenomenology of visual attention, but rather the effects of attention on visual phenomenology. That is, the phenomenology at issue is not a feature of attention but the result of attention. I then consider proposals that answer question (2) affirmatively in Section 4.4, where the phenomenology of visual attention is explicated in terms of uniform changes in visual content. I argue that these proposals fail to identify a uniform phenomenology of attention. Section 4.5 discusses two recent accounts of the phenomenology of attention that do not appeal to perceptual content. Section 4.6 then shifts to a popular philosophical theory of consciousness that explains sensory phenomenology by appeal to the representational content of perception, namely *representationalism*. Various counterexamples to representationalism draw on attention's effects on perceptual experience, but empirical work can be used to set the objections aside. Section 4.7 concludes by considering an interesting new challenge to representationalism by Ned Block that draws on empirical work on attention. I suggest that his argument raises questions about an empirical version of representationalism about consciousness, one that ties consciousness to information processed by neurons.

4.2 Unconscious attention

If consciousness is of attention's essence, then there is no attention without consciousness and hence, no unconscious attention. Admittedly, it is hard to think of mundane cases where attention is unconscious. Every deployment of attention of which you are aware is seemingly tied to conscious experience. How could it be otherwise? Still, what is the evidence or basis for endorsing the strong claim that attention is essentially conscious? Certainly, many endorse this claim, but why? Consider this possibility: confidence in James' essence claim is due to an induction from the fact that attention is recessarily connected to consciousness to the generalization that attention is necessarily confident that all swans are white if all the many swans one has seen are white, but a single black swan would spoil the generalization. The question then is whether there are any black swans where attention is concerned, namely cases of unconscious attention?²

How might one demonstrate that there are cases of unconscious attention? To do so, we need a way to track attention where consciousness is absent. Given that this is a contentious area, one way to find common ground is to go back to the shared assumption discussed in earlier chapters, namely the empirical sufficient condition as applied in standard paradigms used to probe attention. With that in mind, there is compelling experimental evidence that attention can be deployed in unconscious vision. Consider work by Robert Kentridge and coworkers (1999) with the blindsight patient GY. Blindsight patients are understood not to have any visual phenomenal experience of part of their visual field due to damage in primary visual area V1, the first area in cortical visual processing.³ For example, damage to the left primary visual area (V1) will lead to a "blind field" in the right "hemifield", i.e., the right half of the visual field (left hemisphere cortical visual areas process the right side of visual space, the contralateral side, while the right hemisphere processes the left side of visual space). This blind field can be of different sizes and can be precisely mapped (on different conceptions of blindness, see Chapter 6, Section 5).

The damage to V1 accounts for the "blind" in "blindsight", i.e., the subject's defective vision, but what about the "sight"? The basic finding is as follows: blindsight patients still have some access to visual information from the blind field. This information is carried in preserved sub-cortical visual pathways via the superior colliculus that bypass V1 and connect to later cortical visual areas such as V2, V3, V4 and MT (see Chapter 2 for a map of some of these regions in the primate brain; see Weiskrantz 1999 for a discussion of these areas in relation to blindsight). Accordingly, this information is in some way available to patients, for when they are forced to guess about aspects of items in their blindfield, their guesses are above chance. For example, they can detect the presence of targets, as well as their features. Blindsight patients have access to relevant visual information even if not through normal phenomenal experience.⁴

The gist of Kentridge et al.'s work is the use of the Posner spatial cueing paradigm with GY (see Chapter 1, Section 6), a standard test for spatial attention. Kentridge et al. showed that GY demonstrated the standard cueing advantages and costs that are indicative of spatial attention in normal subjects where the relevant space is in the blinded portion of the visual field. In other words, for a normal subject, such effects would be evidence for spatial attention to a location. If one accepts the Posner paradigm as a test for spatial attention, there is prima facie evidence that GY can deploy spatial attention to X without visual phenomenal experience of X. If the empirical sufficient condition is appropriate in spatial cueing with normal patients, then it is appropriate in the case of GY. One should conclude that GY exhibits attention without consciousness. GY selects the target to inform his response, but there is nothing to support the phenomenal notions of focalization, concentration, or spotlighting in his blind field. If you are concerned that these are experiments with a neuropsychological patient, note that similar experiments have also been conducted with neurologically intact subjects (Jiang et al. 2006; Kentridge, Nijboer, and Heywood 2008; Norman, Heywood, and Kentridge 2013). Accordingly, there is empirical evidence that attention can be deployed to unconsciously seen targets and, in that sense, attention can be unconscious.

I suggest that current empirical evidence points to a negative answer to question (1).⁵ My case draws on a widely accepted sufficient condition as a way to leverage empirical data to answer the question whether attention can be unconscious. If one is inclined to deploy the empirical sufficient condition with normal subjects, then there is good reason to apply the condition to blindsight subjects who evince the same behavioral capacities. Now a

proponent of the phenomenal conception of attention who answers (1) affirmatively can take issue with the sufficient condition. For if attention is essentially conscious, then selection for task by a blindsighter will not entail attention. Fair enough, but at this point, one needs a positive argument for an affirmative answer to (1). It is not enough to simply insist, say from introspection, that attention is essentially conscious. How could introspection establish that? Have proponents of the phenomenal conception used introspection to ascertain the complete absence of unconscious forms of attention? Does introspection give them insight into the nature of attention? I am doubtful. Perhaps my own powers of introspection are insufficiently discerning, but notice that an affirmative answer to (1) endorses a striking segregation of consciousness in the mental domain, for over the past century, theorists have come to accept that mental states have conscious and unconscious forms, whether perception, thought, memory, emotion, or desire. An affirmative answer to (1) holds that attention is different from other mental states in that it must be conscious. Such a claim requires a positive argument. In the absence of a concrete proposal by the phenomenal conception, there are reasons to doubt an affirmative answer to (1).

To counter the empirical evidence against (1), a proponent of the phenomenal conception cannot simply insist that attention is essentially conscious. It is high time to present a detailed characterization of the phenomenal conception as a first step in countering the empirical evidence and the empirical sufficient condition. The next three sections try to do this on behalf of the phenomenal conception but with mixed results.

4.3 Attentional effects on visual experience: empirical work

Earlier, I emphasized the idea that conscious attention should involve a uniform phenomenal character, a distinctive attentional phenomenology. The alternative is that conscious attention involves disparate phenomenology, different ways of being conscious. Let us consider this alternative, namely an affirmative answer to (1) but a negative answer to (2). The problem is that disparate phenomenology undercuts the phenomenal conception. If one thinks that attention is essentially conscious, the endorsement of disparate attentional phenomenology seems initially odd, for how can something be essentially a distinctive mode of consciousness, yet at the same time involve disparate phenomenology? Well, one might reply,

we allow that conscious vision has disparate phenomenology. For example, sometimes visual phenomenology involves the appearance of chromatic colors and sometimes it does not, involving instead the appearance of only achromatic colors. Still, might conscious vision have a uniform phenomenology in that it always involves some color phenomenology? Were one's experience not to have color phenomenology, one would not be having visual experience. But set that worry aside. The idea is that, as with conscious vision, so with conscious attention: one can allow for a distinctive attentional phenomenology even if the phenomenology is disparate. But this raises the issue of what attentional phenomenology comes to. This section looks at empirical work that highlights what attentional phenomenology might be in the visual domain. Drawing on work by Marisa Carrasco and co-workers, we shall see that in vision, attention does seem to have varied effects on visual content, say alterations to visual representations of perceived size, contrast, or saturation of color. While this suggests that attention is tied to disparate phenomenology, the phenomenology is not a feature of attention, but rather the effect of attention. That is, the empirical work shows how attention affects the phenomenology of visual experience in disparate ways. As such, it does not reveal the phenomenology of attention, different ways it is like to be in a state of attention.

4.3.1 Contrast

At the beginning of modern psychology, Gustav Fechner and William James disagreed on whether attention affects how things perceptually appear. Fechner claimed that "a gray paper appears no lighter, the pendulum-beat of a clock no louder, no matter how much we increase the strain of our attention upon them" (quoted in James 1890 p. 426).⁶ On the contrary, William James writes:

Every artist knows how he can make a scene before his eyes appear warmer or colder in color, according to the way he sets his attention. If for warm, he soon begins to *see* the red color start out of everything; if for cold, the blue. Similarly in listening for certain notes in a chord, or overtones in a musical sound, the one we attend to sounds probably a little more loud as well as more emphatic than it did before.

(James op. cit. p. 425)

How can two noted psychologists disagree on the phenomenology of attention, something one might think to be clear on careful introspection? Who would you side with?

One worry is that introspective reports will be biased by theoretical commitments. Marisa Carrasco identified a clever way to reduce this potential bias by probing subjects' conscious awareness indirectly. Specifically, Carrasco probed conscious awareness by requiring subjects to make assessments about how things visually appear in order to explicitly report on a different feature. Subjects never explicitly report on their conscious experience in the task, but experimenters can make inferences about aspects of their experience in light of their performance. Let us examine Carrasco's first study, focusing on the perception of contrast.

The perception of contrast is a fundamental feature of visual experience. In a natural scene, think of differing levels of luminance as in the shadowy pattern cast by a picket fence. It will suffice to think of contrast using the following diagram, essentially a sequence of black and white lines where the transition in luminance between them is gradual and not abrupt:



Figure 4.1 Two contrast grating patterns of different spatial frequencies, here measured in terms of cycles per visual degree: (A) one cycle per visual degree; (B) two cycles per visual degree. As demonstrated, each pattern can be represented by a sinusoidal wave. Figure from Webvision, http://webvision.med.utah.edu/ and is courtesy of Michael Kalloniatis and used with his permission.

This "grating" can be represented in terms of a sinusoidal wave function that represents changes in luminance across the display as well as the spatial frequency of the display. The contrast of such gratings can be calculated in different ways, but a standard way for simple periodic patterns is the Michelson Contrast, C_m , which is the difference between the maximum and minimum luminance in the pattern divided by their sum.⁷ This yields a value between 0 and 1 and is often expressed as a percentage. A standard psychological stimulus illustrating periodic changes in contrast is the *Gabor* patch, a visual stimulus that is constructed by multiplying a sinusoidal wave function with a Gaussian function, leading to a pattern that decreases in amplitude from its "center". Here is a set of stimuli of the sort used by Carrasco with their contrast values:



Figure 4.2 A set of Gabor patches at different levels of contrast (panels (A) and (B), the former including low contrast stimuli; the latter, high contrast). When fixating on one of the two dots, attending to the Gabor patch on the left seems to make it appear to be of the same contrast as the patch on the right. Reprinted by permission from Macmillan Publishers Ltd: Carrasco et al. (2004) "Attention Alters Appearance." Nature 7: 308–313.

Carrasco and colleagues had previously shown that attention improves a subject's sensitivity to contrast and argued that this was due to attention's enhancing visual representation of contrast (Carrasco, Penpeci-Talgar, and Eckstein 2000). In subsequent work, Carrasco and coworkers investigated whether this signal enhancement was reflected in visual experience, namely in the visual appearance of contrast (Carrasco, Ling, and Read 2004). They did not directly assess subjects' reports of experience; rather, subjects reported on the orientation of a Gabor patch in light of the patch's apparent contrast. Subjects were presented with two Gabor patches of different contrasts and reported the orientation of the patch that appeared highest in contrast (the patches were tilted 45° either left or right). The experiment begins with subjects fixating on a central dot. This is followed by a cue either at



Figure 4.3 Subjects fixated on a central fixation point for 0.5 seconds. A cue was presented for 67 milliseconds (ms), either at the fixation point (neutral cue), or at a location near where one of the Gabor patches would appear (peripheral cue). After an interstimulus interval (ISI) of 53 ms, the two Gabor patches would appear for 40 ms. Subjects were required to respond within one second, reporting on the orientation of the Gabor patch that appeared to them of highest contrast. Reprinted by permission from Macmillan Publishers Ltd: Carrasco et al. (2004) "Attention Alters Appearance." Nature 7: 308–313.

the point of fixation (neutral cue) or at the periphery near one of the targets (peripheral cue). After an interstimulus interval of 53 milliseconds (ms) between cue and stimuli, two Gabor patches appear equidistant from, and on the left and right of, the fixation point for 40 ms. One of the patches is the *standard*, kept at a specific contrast; the other, the test patch, varies in contrast. Subjects were given 1000 ms to report the orientation of the Gabor patch that appears highest in contrast (Figure 4.3). The experimental design uses peripheral cues to tap into bottom-up or endogenous attention, a form that the experimenters note peaks at about 120 ms after the cue and largely decays by 250 ms (see Figure 1.4). This explains the short presentation time of the stimuli, which attempts to take full advantage of bottom-up attention.

A key feature of the experiment is that subjects did not explicitly report which patch appeared highest in contrast. Rather, their reports about orientation depended on how the contrast of the patches appeared to them (for a discussion of the logic of the experiment, see Carrasco 2006). Carrasco discovered that attention appeared to make the Gabor patch appear to be higher in contrast than it in fact is.⁸ In experiments where the standard Gabor patch was of 6% contrast, attention appeared to boost apparent contrast by 2.5%. Thus, an attended 3.5% Gabor patch would appear to be the same as a 6% contrast patch. With a higher contrast standard of 22%, attention appeared to boost apparent contrast by 6% (Figure 4.2). In the original experiment, the cue brings bottom-up attention on line (Carrasco refers to this as exogenous or transient attention), and later work demonstrated similar effects with top-down attention (Liu, Abrams, and Carrasco 2009).⁹

This is a striking result. While there have been some queries raised about methodology and alternative explanations, Carrasco has responded to these in detail, and I shall take on her interpretation.¹⁰ There is also work showing that the contrast effects can be achieved crossmodally where attention is cued in an auditory fashion, affecting the visual experience of contrast (Störmer, McDonald, and Hillyard 2009). The upshot is that there is compelling empirical evidence that attention affects the visual appearance of contrast.¹¹

4.3.2 Gap size

Using the same experimental design as deployed in the contrast work, Carrasco and colleagues have investigated the effect of bottom-up attention on discrimination of gap size using Landolt squares (squares that have a small gap in one of the sides).¹² In earlier psychophysical work, Carrasco, Williams and Yeshurun (2002) found that discrimination performance improved in the peripheral versus neutral cue conditions on both reaction time and accuracy, suggesting that when attention is drawn to the target location, processing is facilitated. Similar results were seen for top-down attention as well (Montagna, Pestilli, and Carrasco 2009). Indeed, Montagna et al. noted that there was also a corresponding cost in discrimination of squares in unattended locations. Are these enhancements in performance reflected in conscious visual experience?

Joetta Gobell and Marisa Carrasco (2005) used a variation of the original paradigm deployed in the contrast work: subjects were presented with two Landolt squares and reported on the orientation of the gap (top or bottom) in the square that had the largest apparent gap. One of the squares, the standard square, had a gap that spanned 0.2 visual degrees while the other square, the test square, had a gap that spanned a range between 0.04 to 0.36 visual degrees (your thumb at arm's length has a width of about 2 visual degrees). The cue was uninformative as to the orientation of the gap (top or bottom) and it directed attention either to the standard square or to the test square.



Figure 4.4 Diagram of the experimental task conducted by Gobell and Carrasco (2005). Subjects fixated on a central dot. A cue, either peripheral (left panel) or neutral (right panel) was presented, followed by an interstimulus interval of 40 ms. The stimuli consisted of two Landolt squares presented equidistant from fixation. Subjects then had to respond in less than two seconds whether the gap in the square with the largest gap was on the top or the bottom of the square. Figure reproduced from Gobell and Carrasco (2005). Reprinted by permission of SAGE Publications.

Three conditions are relevant: (1) attention to the test square; (2) attention to the standard square; (3) attention to neither (neutral cue). In all cases, the standard square gap is 0.2 visual degrees. Gobell and Carrasco found that in the attention to the test square condition, the test gap was judged to be the same size as the standard (0.2 visual degrees) when its gap was 0.18 visual degrees; in the attention to the standard square condition, the test gap was judged to be the same size as the standard square condition, the test gap was judged to be the same size as the standard square condition, the test gap was judged to be the same size as the standard square condition, the test gap was judged to be the same size as the standard gap when the former was 0.23 visual degrees; in the neutral condition, judgments of same gap size occurred when the two gaps were roughly of the same size. Gobell and Carrasco conclude that attention "increases apparent gap size in a Landolt-square acuity task" (644).¹³

4.3.3 Color

Carrasco has also examined whether attention affects the perception of the three dimensions of color: namely, saturation, hue, and intensity. The studies of contrast already suggest that attention can affect intensity, but what of hue and saturation? Think of the dimension of hue as, say, blue versus red and the dimension of saturation as deep blue versus pale blue (recall James's earlier remark on seeing warmth and cold). Stuart Fuller and Carrasco (2006) adapted the previous paradigms to examine the effects of attention on the perception of hue and saturation. As in the experiments

discussed above, subjects were required to make orientation judgments regarding two stimuli on the basis of their relative color saturation or hue. In the saturation dimension, subjects were asked to judge the orientation of patches (tilted left or right) that were "redder, greener, or bluer". In the hue dimension, subjects were asked to judge the orientation of patches that looked more blue or more purple. Fuller and Carrasco report that attention affects the appearance of saturation but not of hue. They explain this difference in light of a difference in the two properties: saturation, like contrast and spatial frequency, can be understood as a magnitude allowing for more or less, while hue cannot. It might be that attention only affects the representation of properties that one can think of in terms of more or less such as size, speed, contrast, and saturation.

4.3.4 Brightness

Finally, Peter Tse (2005) has presented an illusion that deploys top-down attention. I raise this example because it is different from the Carrasco effects that require very short presentation times of the central stimuli. The illusion is quite striking and seems to be induced by how one attends to the display. One's experience changes by shifting attention at will, though further work needs to be done to understand how this illusion is induced by attention (Tse emphasizes this as well).¹⁴

In the right panel, fixate on any of the dots and while maintaining fixation, shift covert attention to any of the three overlapping disks. Introspect on your visual phenomenology. Anything? Now do the same thing for the



Figure 4.5 Illusion due to Peter Tse. From P. Tse (2005) "Voluntary attention modulates the brightness of overlapping transparent surfaces." Vision Research 45: 1095–8 with permission of Elsevier.

left panel and again introspect on your visual phenomenology. Is there a difference between the two experiences? Subjects reported that in the left but not right panel, attending to a disk makes the disk appear darker. This seems to work more strongly in the left figure, but not the right (see Tse op. cit. for an explanation). The point is that attention seems to affect phenomenology, in this case decreasing apparent brightness of the attended disk (or increasing the relative contrast between the disks). The effect is striking, and the difference between the two displays seems to point to a clear effect on visual experience of brightness that tracks attention.

4.3.5 Disparate attentional phenomenology

I have now drawn on Carrasco's work and an example from Tse to suggest different phenomenal consequences of attention in vision. This includes visual phenomenology in respect of perceived size, contrast, brightness, and color. If the work is correct, then the deployment of attention, both top-down and bottom-up, can lead to different effects in visual phenomenology. These results, however, do not support a phenomenal conception of attention that holds that attention has disparate phenomenology. Recall that the phenomenal conception of attention takes attention to be a distinctive mode of consciousness and so to have a distinctive phenomenal character. The phenomenal conception construes the phenomenology to be a property of the state of attention, just as visual phenomenology is a property of visual experience. It is, after all, a conception of the phenomenology of attention. Yet the empirical data indicate that the relevant phenomenology concerns changes in visual phenomenology that is caused by attention. This suggests a weaker conception of conscious attention: attention is conscious in that it brings about (possibly distinctive) phenomenology in other mental states but attention does not have its own phenomenal character, a phenomenal feature of the state of attention. This weaker conception points to a disanalogy with vision, for where a visual experience might have disparate phenomenology (say for color, texture, shape, etc.), that phenomenology is a property of the experience and not a property brought about in other states. Attention is different, for the disparate phenomenology it is tied to is brought about in other states such as visual experience. It is not clear then why one should speak of the resulting phenomenology as specifically attentional since it involves changes in visual phenomenology such as alterations in apparent contrast, size, and saturation. Couldn't such changes occur in visual experience without attention?¹⁵ Moreover, it is hard to see that such phenomenology is of attention's essence or that it is a necessary aspect of what it is like to attend. The appeal to empirical work to explain attentional phenomenology does not, then, aid a phenomenal conception of attention. Attention is not revealed to be a distinctive mode of consciousness, but rather, at best, to affect other forms of consciousness in disparate ways. As we do not as yet have good reason to think that the phenomenology of attention is disparate, I suggest that we pursue a characterization of that phenomenology as being uniform.

4.4 Attention and visual experience

I have argued that the best approach for the phenomenal conception of attention would be to identify a uniform phenomenology of attention. Until that conception does so, its assertion that attention must be conscious is empty, for one does not have an idea of what conscious attention might be. The previous empirical results suggest that attention can have various and disparate effects on visual experience, but those results do not sit well with the phenomenal conception. Still, might there be a way of unifying these disparate features so that there is a uniform phenomenology associated with attention, even if this phenomenology concerns another state like visual experience? Consider the following claim of a phenomenal conception:

When conscious attention is deployed in conscious visual experience, there is an "attentional phenomenology" of type C that necessarily qualifies visual experience.

Think of this relation between attention and C as holding of necessity in some sense, and thus that every case of conscious attention in vision is accompanied by a uniform phenomenology of type C. The presence of C suffices for conscious visual attention. Now, even if one can identify a relevant C for vision, this still raises the question of whether one might find an equivalent C for each of the other sensory modalities, such as audition and touch, and indeed, for non-sensory modalities where attention might have a phenomenal upshot, such as emotion and thought. If we impute disparate phenomenology across each modality, then we must confront the plausible possibility, in light of empirical work on vision, that the phenomenology is the result of attention in each case and not indicative of a phenomenology of attention. But set that worry aside. Can we even

identify a relevant *C* in vision? Does *C* always track attention? In this section, I draw on Carrasco's empirical work to formulate an answer in terms of perceptual representational content. That is, I draw on what is represented in perception to provide proposals for *C*.

Carrasco's work suggests that attention affects conscious visual experience across a wide range of visible parameters. If correct, her results show that the subject's state of attention is a determinant of some of the contents of the subject's visual experience: the experience of contrast, of saturation, and of gap size. As these are visual features that pervade normal visual experience, attention's effects will also thereby be pervasive. Change the subject's state of attention and there will be correlated changes in the subject's experience of many visual parameters. It is an empirical question how broadly distributed the effects of attention are, for as noted earlier, the experience of hue might not be modulated by attention. The empirical results are striking, but it is not clear that they demonstrate a uniform phenomenology of attention in the sense of a uniform effect on visual consciousness. Indeed, attention-modulated shifts in visible gap size, changes in visible contrast, and increases in apparent saturation seem quite disparate phenomena. Might there be a way to unify these changes in terms of identifying a common phenomenal characterization?

James Stazicker (2011) and Bence Nanay (2010) have both emphasized the idea that attention increases the determinacy of visual representations. This suggests one way to flesh out a phenomenal conception of attention, namely, by emphasizing that conscious attention is so called because it is always tied to increases in the determinacy of visual representations.¹⁶ Why are representations of relevance to phenomenology? When one speaks of perceptual phenomenology, one often has in mind perceptual content in the sense of perceptual representations of features of the experienced world. In later sections, I shall discuss representationalism, a theory that explains the phenomenal character of experience by appeal to the representational content of experience. This theory aside, it does seem that for many aspects of perceptual phenomenology, there is an associated perceptual content. So, the phenomenology of seeing a Gabor patch to be of high contrast is tied to a visual content representing that patch to be of high contrast; the phenomenology of seeing a color to be of greater saturation than another is tied to a visual content representing this difference in saturation. The appeal to determinacy appeals to a feature of visual content to explain the phenomenology of conscious attention. One's visual experience of an X when one attends to it is more determinate in respect of the

experience's representing X's features than when one experiences X without attending to it.

Stazicker and Nanay explain the idea of increases in determinacy in terms of the determinable/determinate distinction, where a determinate of a determinable can be understood as a more specific way of being that determinable. Thus, scarlet is a determinate of the determinable red with the latter itself being a determinate of the determinable color. Scarlet is a (more specific) way of being red and red is a (more specific) way of being colored. Another way to think about an increase in determinacy is in terms of a decrease in uncertainty about the feature in question. Being between 0.8–0.9 meters is a more determinate way of being less than 1.0 meters, and it reduces uncertainty about the length of the line. There is more uncertainty about the exact length if you provide a larger range as opposed to a smaller range. The phenomenal conception of attention then can associate conscious attention with increases in perceptual determinacy.

Todd Ganson and Ben Bronner (2012) point out, however, that some of the effects identified in Carrasco's work do not involve increases in determinacy. For example, if attention's effect on contrast is to be understood as boosting absolute contrast level, say from 3.5% to 6% contrast, then this is not an increase in determinacy as would be the case if attention shifted contrast representations from 2-7% contrast to 3-6% and thus narrowed the represented range (the same point holds for relative contrast between the two patches, say a shift from unequal to equal contrast).¹⁷ Similarly, changes in experienced saturation might be due to an increase in the level of saturation as experienced. Thus, even where attention affects visual content, it need not be tied to an increase in determinacy. So conscious attention can be disconnected from increases in determinacy. If this is correct, then an increase in determinacy is not a viable proposal for *C*.

Elsewhere, I (Wu 2011c) and, independently, Sebastian Watzl (2011) have argued against various attempts to appeal to some aspect of perceptual content to characterize C. To give a sense of how the argument goes, consider two illustrative proposals for C in the visual domain: attentional phenomenology as tied to an apparent boost in contrast, or as tied to an apparent increase in spatial resolution. One might claim that when one attends to a Gabor patch, the relevant phenomenal upshot is that contrast appears greater than it did earlier when it was not being (focally) attended (see the notion of diachronic salience in Section 4.5). The phenomenology of attention is, then, the increase in the level of apparent contrast of a patch over time, between its being unattended versus its being attended. Yet, in

principle, one might be able to manipulate the stimuli or even the neurons processing that portion of the visual field such that the apparent contrast of the patch decreases over time. If so, then the change in apparent contrast will be in the opposite direction of the proposal that attention boosts apparent contrast, yet it seems like one can attend to objects that slowly grow dimmer and whose level of contrast decreases with time, e.g., as one attends to an object in the diminishing light as the sun sets.

One might also claim that attention always increases spatial resolution. Yet, the worsening of visual spatial resolution seems to be compatible with attending to an object that becomes less spatially resolved. Think of attending to letters on a screen through a lens as an optometrist adjusts various lenses in the search for one's correct eyeglass prescription. One can attend to the letters even though the object seems blurrier over time. Indeed, optometrists rely on this capacity for attention while they make adjustments that change the apparent spatial resolution of visual experience in order to generate a correct prescription. In general, attention to a dimming object in the waning light or attending to an ever more blurry letter in an optometrist office is not generally characterized as unconscious attention. So, it seems that the two proposals for C can be decoupled from conscious attention.

The strategy is to take any proposal for C based on perceptual content and show that it is possible to have attention without C. As I have not exhaustively considered all available options, it would be premature to draw the strong conclusion that no uniform phenomenology as tied to perceptual content will be present in each and every case of visual attention. Nevertheless, I have canvassed the cases that are initially most plausible, in that they are grounded in Carrasco's work (for arguments against more complicated proposals for C, see Wu 2011c). If there are to be viable proposals for a uniform phenomenology tied to attention, then proposals from proponents of the phenomenal conception must be given. Until then, I suggest a tentative conclusion: for all that has been proposed thus far or that are reasonable extrapolations from those proposals - namely, increases in determinacy or boosting of perceptually represented magnitudes like contrast attention does not necessitate a uniform phenomenology tied to some aspect of perceptual content. Moreover, the effects of attention on perceptual content remain a motley, non-uniform sort. This conclusion can allow that attention has a typical phenomenology in that, for example, when attention is directed at contrast, the effects that Carrasco suggests might generally obtain, e.g., a boost in apparent contrast or an increase in the determinacy of spatial representations. That such effects are typical can also explain why

theorists gravitate towards specific metaphors to describe attention. A typical effect, however, is not an essential effect.

4.5 Salience and structure in the phenomenology of attention

In the last section, I argued that there is as yet no viable proposal that identifies a distinctive essential phenomenology of attention. Coupled with the empirical evidence for unconscious attention, I believe (1) receives a negative answer: attention is not essentially conscious. But what of an affirmative answer to question (2): attention, when conscious, has a distinctive and uniform phenomenology? The empirical evidence does not appear to identify a uniform phenomenal effect of attention, but in this section, I canvass two proposals that identify a specific attentional phenomenology. Both proposals are unified in that they do not appeal to perceptual content to explicate the phenomenology of attention. Moreover, both are general, not tied to vision. The proposals are:

Phenomenal Salience: the phenomenology of attention is the rendering of the attended object as phenomenally salient.

Phenomenal Structuring: the phenomenology of attention is the structuring of the conscious field around the attended object.

Let us begin with phenomenal salience (Wu 2011c). Of course, "phenomenal salience" is just a label for the putative uniform phenomenology of attention, something we have yet to pin down. It is thus to be separated from the use of "salience" in psychological contexts where it often refers to a property of a stimulus that draws attention to it (e.g., a red singleton among green objects is salient and "pops out"). The intuitive starting point is that, in perception, attention highlights or phenomenally prioritizes a specific item. Thus, pick two perceptual objects in each modality and switch attention from one to the other: hold two objects in each hand, and switch tactile attention between them; identify two concurrent sounds, and switch auditory attention between them; look at two objects before you, and keeping the eyes fixed, switch attention between them. Really try it! Does the attended object in some way seem more salient than the other? If no, then skip to the end of this section! If yes, then that is the *explanandum*.

One version of the phenomenal conception might simply assert that phenomenal salience refers to recognizable phenomenology that is available to subjects on introspecting their episodes of attention. This phenomenology, moreover, is sui generis or something primitive (for the suggestion though not necessarily the endorsement, see Speaks 2010; Chalmers 2004). The best way to counter this position is to show that phenomenal salience can be explained by appeal to components that are already at hand without needing to hold that attentional phenomenology is a primitive, irreducible feature of conscious experience. The explanation of phenomenal salience that I have offered elsewhere attempts to do just that (Wu 2011c).

Phenomenal salience has two temporal dimensions:

Diachronic Phenomenal Salience: The salience of an object at some time relative to that object at an earlier time.

Synchronic Phenomenal Salience: The salience of an object at some time relative to other objects at the same time.

Thus, when you shift attention between two Gabor patches, you toggle the diachronic phenomenal salience of a single patch over time as you switch attention to and fro. When attention lands on a patch, that patch is more phenomenally salient than it was at an earlier time when attention was directed elsewhere. At the same time, when you attend to one of the patches, you render it synchronically more phenomenally salient relative to the other patch. Of course, what phenomenal salience amounts to is the open question. The point of the previous section was to emphasize that there is as yet no adequate account of a uniform attentional phenomenology. Let me now offer such an account. I have proposed (Wu, op. cit.) that phenomenal salience is a product of special cases where subjects explicitly reflect on the phenomenology of attention. Normally, when one shifts attention to and fro, there is often no accompanying phenomenal salience. Mundane visual experience does not seem to involve a movable spotlight, or else one is confusing the effects of foveation for an attentional spotlight.¹⁸ Think of all the eye movements you have just made in the past few minutes. Was that like a highlighting or shifting spotlight? Yet in reflecting on the phenomenology of attention, as we have been doing, we render the target of attention salient in thought. If thinking is a form of cognitive attention, then the idea is that in reflective cases where phenomenal salience seems, well, salient, a perceived object is the target of perceptual and cognitive attention, i.e., thought. We reflect on how attention presents that object rather than others. More generally, we are in a sense aware that we are attending to that object. In this way, we make the target phenomenally salient not just because we perceptually attend to it, but also because we think about it. The suggestion, then, is that phenomenal salience is a function of the addition of cognitive attention's locking onto perceptually attended items. The additional selectivity of thought explains phenomenal salience. If correct, then perceptual phenomenology does not explain phenomenal salience. Phenomenal salience is all in reflective thought. Specifically, the proposal is that conscious attention, in terms of phenomenal salience, is a property of perception-dependent *demonstrative* thought, a topic discussed in Chapter 7. As such, phenomenal salience is a product of special situations when we reflect on attention. Since we are not always in this way reflective, phenomenal salience does not always track attention.

Sebastian Watzl takes a different approach. On his account, the phenomenology of attention is to be understood in terms of structuring the conscious field. Watzl expresses this as follows:

consciously attending to something consists in the conscious mental process of structuring one's stream of consciousness so that some parts of it are more central than others.

(Watzl 2011, 158)

The phenomenology of attention is tied to a structure imposed on the conscious field by attention that yields a center/periphery structure. Consider our experience of Figure 4.2 and reflect on the difference between three experiences: E1 where you are attending to the 22% Gabor patch, G, on the bottom; E2 where you are attending to the fixation point to the right of G; and, E3 where, if possible, you are not visually attending to the figure at all, but, while looking at the figure, attending to a thought. Attention then renders G more central and the fixation point more peripheral in E1; it switches centrality and peripherality in E2; and it leaves a flat structure, at least with respect to visual experience, in E3 (i.e., there is no center or periphery in the phenomenal structural sense in the visual field). This center-periphery structure obtains as well in considering conscious experience as a whole. When one focuses attention in vision, one renders the attended visual object as central, but other visual objects, auditory objects, tactile objects, and so on are more peripheral.¹⁹ The phenomenology of attention here is not captured in terms of representational content, but in terms of a structural feature of experience, i.e., the relation between the experience's parts.

There is something very compelling about Watzl's characterization of the phenomenology of attention, but it is also in a way elusive. In discussing phenomenal salience, I suggested that the phenomenology of attention is a product of very special cases where one reflects on attention and its targets. In reflecting in this way, one focuses on the targets, although in a way that need not involve a change in the phenomenology of experience per se. Rather, focalization, concentration and now, centering, results from the special status of the target of attention, namely, that it feeds in a direct way into one's actions such as reflective thought. Watzl might be right that there is something like a center-periphery structure, but I claim that this is a reflection of special cases. Alternatively, Watzl can claim that I have simply missed a common structural feature of all perceptual experiences where attention is differentially deployed.²⁰ At this point, the debate mirrors something like the exchange between James and Fechner, a difference in basic intuitions about the phenomenology of attention. The challenge then is how to resolve the impasse when one hits rock-bottom disagreements about how consciousness seems to be. Watzl sides with James in finding attentional phenomenology to be widely distributed across experiences; I, on the other hand, side with Fechner in finding less attentional phenomenology.

This chapter began with a phenomenal conception of attention that takes attention as essentially conscious, a distinctive mode of being conscious. Against this, I adduced empirical evidence for unconscious attention drawing on the empirical sufficient condition. Those who endorse a phenomenal conception have reason to deny that condition, but the challenge for them is to characterize the idea of attention as a distinctive mode of consciousness in that it has or yields a distinctive phenomenology. To move beyond metaphorical characterizations like the spotlight, I have drawn on Carrasco's work, but that work seems to demonstrate that rather than having its own distinctive phenomenology, attention affects the phenomenology of other states and in quite disparate ways. One possible way to impose uniformity on attentional phenomenology is to appeal to higher-order characterizations such as increase in determinacy, but this does not seem to be found in all forms of attentional phenomenology. Rather, if there is a uniform phenomenology to attention, it might be limited to specific cases and explained by either a cognitive component (phenomenal salience) or a higher-order, nonrepresentational structure of experience (phenomenal structuring). But at that point, the phenomenology might seem elusive, or so it seems to me. The current state of the dialectic is tilted towards a negative answer to (1) but a potentially positive, if limited answer to (2). A negative answer to (1), however, undercuts a phenomenal conception of attention, even if attention often has a phenomenal (possibly uniform) upshot. Accordingly, a phenomenal conception of attention does not provide an answer to the metaphysical question, as not all forms of attention are conscious.

4.6 Attention and representationalism

Let us shift gears and consider another link between attention and perceptual phenomenology. The Carrasco results are relevant to philosophical theories of perceptual consciousness. A currently influential account of perceptual consciousness is *representationalism* which aims to explain the phenomenal character of experience in terms of its representational content. The appeal of this approach is that it accounts for something many find mysterious, namely, phenomenal consciousness, in terms of a facet of perception that many find explicable, namely representational content (for an overview of representationalism, see Chalmers 2004). That is, one can explain what perceptual experience is like by appeal to what perceptual experience represents. In its strongest version, representationalism claims that phenomenal character just is representational content that meets certain conditions (Tye 1995). A common form of representationalism, however, claims that phenomenal character supervenes on representational content:

Supervenience of Phenomenal Character on Content: Necessarily, if two experiences differ in phenomenal character, then they differ in representational content.²¹

There are different versions of this claim, but I shall focus on perceptual experiences within the same sensory modality, and hence on a claim about intramodal supervenience.²² Accordingly, visual phenomenal character supervenes on visual representational content such that two visual experiences that differ in phenomenal character also differ in what they represent. The idea of representational content is tied to the notion of an *accuracy condition*. That is, an experience with representational content p is accurate or veridical if and only if p. So, a visual experience that there is a blue car passing before one is veridical if and only if there is a blue car passing before one. The phenomenology of that experience is then explained by its content, namely, that a blue car is passing before one. It is also worth noting a type of representationalism that is often in play in the neuroscience of consciousness in the Content Realization Principle (CRP):
(CRP) There is a necessary correlation between the content of consciousness and the information carried by the neural realizers of consciousness.

Here, conscious visual content supervenes on information processed in the visual system. Many cognitive scientists of consciousness endorse a form of representationalism via CRP. It is supervenience that is the target of many putative counterexamples to representationalism.

Let's note an alternative to representationalism, namely relationalism.²³ Relationalism denies that perception is in any substantive sense representational, namely, as having a content that determines an accuracy condition, and thus, denies that phenomenal character supervenes on representational content. Rather, perceptual consciousness involves in part an awareness relation to what one is aware of, and the features of the objects of awareness determine the phenomenal character of the experience. For some relationalists, the objects of awareness are mental entities such as sense-data; for others, the objects of awareness are physical objects. In the latter case, the phenomenology of visual experience will be determined by the visible properties of the objects of visual awareness (Campbell 2002). A visual experience of a red rose has the phenomenal character of the perceived object determines the phenomenal character of the experience.

David Chalmers (2004) has suggested that attention might yield the most plausible counterexamples to the representationalist thesis of supervenience.²⁴ Let us now quickly survey a few salient cases that have been offered. The essential strategy is to note that when attention is deployed, there is a phenomenal consequence in experience that cannot be explained by appeal to representational content. Each of these examples looks initially very compelling, but the challenge they pose to representationalism can be met once Carrasco's results are in view. An adequate representationalist response to counterexamples, however, doesn't just adduce changes in content willy-nilly with every change in phenomenal character. Rather, plausible responses must hew to the following claim:

Representational Intelligibility: Necessarily, for any experiences E1 and E2, if there is a phenomenal difference between them, then there is a representational difference between them that makes the phenomenal difference *intelligible* (i.e., bears some explanatory relation to it).

(Wu 2011c, 101)

That is, the representational difference must make sense of the phenomenal difference. Exactly what this intelligibility or explanatory relation comes to will be left open here. It is a vexed issue, tied to the difficult question of what would count as explaining consciousness, but it will be enough that in each representationalist response to be considered, the adduced content is plausibly connected to the phenomenology.

Bernhard Nickel (2007) has argued that representationalism fails given the following nine tiles figure (Figure 4.6):



Figure 4.6 As discussed in Nickel (2007) using the number scheme on the right in panel (B), first attend to boxes 2, 4, 6, and 8, then shift attention to boxes 1, 3, 5, 7, and 9. Notice the phenomenal difference?

Nickel notes that one can look at the figure in different ways, say focusing on tiles 1, 3, 5, 7, and 9 or focusing on tiles 2, 4, 6, and 8. Each experience where one groups the tiles in different ways seems to have a different visual phenomenal character. Nickel argues, however, that both experiences have the same visual content. In presenting his case, he canvasses and plausibly rules out a set of options that do not appeal to attention. With Carrasco's result in view, however, one can provide a response. In attending to a set of squares, attention boosts the apparent contrast of the scene, with attended squares seemingly of greater luminance than unattended squares. This response extrapolates from Carrasco's work on contrast, which implicates attentional affects on apparent luminance, but it seems a plausible reply.²⁵

Chalmers's example concerns two red pinpoint lights against a black background, where one shifts covert attention from one light to the other. The basic effect can be achieved in the following figure:

Figure 4.7 Focus on the central cross and then shift attention from the left circle to the right circle. Notice the phenomenal difference?

Fixate on the central cross and cast covert attention first to the left dot, then to the right dot. Both experiences differ in phenomenology due to attention, but what is the difference in content? Again, if the effect is visual, one can appeal to determinacy, contrast, or saturation as potential correlates of the effects of attention. For example, when attending to a specific dot, its contrast might seem higher than that of the other (i.e., their contrast relative to their immediate surrounding area). Given these examples, I offer the following conjecture: for any case where the deployment of (covert) attention yields changes in visual phenomenology, the relevant changes track changes in representational content along the lines indicated by Carrasco's experiments. If so, a certain class of challenges to representationalism based on attention can be met.

4.7 Attention and verdicality

Ned Block (2010) has presented a novel argument against representationalism, drawing on empirical work on attention (he also uses the same materials to argue against relationalism). The problem he raises is, in the first instance, a familiar one: he identifies two experiences that involve different phenomenology without different content. What is different is how he argues that the contents are not different, specifically that there is no way to specify the difference. Representationalists respond to putative counterexamples to supervenience by uncovering differences in content. Block's argument attempts to undercut this strategy.

Look at Figure 4.2, focusing on the fixation point to the right of the 22% Gabor patch. In experience E_A , maintain fixation but attend to the 22% Gabor patch; in E_B , maintain fixation and attend to the fixation point. Carrasco's work shows that in E_A , the two Gabor patches flanking the fixation point appear of the same contrast; in E_B , they do not. That is, you experience their relative contrast. Block focuses instead on the experiences of the absolute contrast of just the 22% Gabor patch, call it G. In one experience, call it E1, attend to G; in the other, E2, do not, but attend elsewhere.²⁶ The centerpiece of Block's argument then is a challenge. Take

experiences E1 and E2 of Gabor patch G as representing G as having a specific contrast. If E1 and E2 differ in how they represent the level of contrast of G, then they can't both be correct. One experience must be wrong and in that sense an illusion: it represents G as having contrast value X when G doesn't have that contrast. But which one is illusory? Block's challenge is that there is no principled way to identify which experience is illusory.

Here is the intuitive version. Attention induces widespread and variegated changes in the activity of visual neurons, yet Carrasco's work shows that attention affects appearances. One might think, then, that attention will have diverse effects across the entire visual field. Since visual content is presumably fixed, at least in part, by neural activity, one can ask what level of attention, and its resulting modulation, determines accurate content. Block's claim is that there is no way to nonarbitrarily determine what this level of attention is. One of the experiences of the Gabor patch (E1 or E2) is illusory (inaccurate), but in not being able to fix content nonarbitrarily, representationalists are not able to provide a principled account of which experience is illusory. This raises a challenge for representationalists, for to account for changes in phenomenal character, they need to identify specific changes in representational content. If attention makes specification of content arbitrary, then it seems that the representationalist will not be able to discharge this explanatory demand.

Block's argument relies on two empirical claims, one deriving from Carrasco's work, and the other from work regarding how the amount of attention varies across the visual field. We can present the argument in detail as follows:

- 1. Assume representationalism: any change in phenomenology is accounted for by changes in content that make the phenomenal change intelligible [for reductio].
- 2. Empirical Claim 1: E1 and E2 differ in phenomenology due to how attention is directed in respect of G (specifically absolute contrast) [Carrasco].
- 3. Attention alters the content of contrast experiences E1 and E2 [1,2].
- Empirical Claim 2: The distribution of amount of visual attention varies across the visual field in complex ways (Datta and DeYoe 2009).²⁷
- 5. Representational contents are precise: Any experience E of Gabor patch G is accurate if and only if the content of E is that G is of 22%

contrast. [Call this accurate content. Other representations of G's contrast are then inaccurate contents.]

- 6. So: There is some amount of attention n directed at G in the visual field where an experience E has accurate content and some value [or set/range of values] m (m \neq n) where E has inaccurate content [from 3–5 and implicitly CRP].
- 7. Any choice of a value for n or m will be arbitrary.
- 8. One can nonarbitrarily assign an accurate/inaccurate content to E only if one can non-arbitrarily identify the appropriate amount of attention for accurate/inaccurate content.
- So, one cannot nonarbitrarily assign an accurate/inaccurate content to E [7,8].
- 10. If one can't assign accurate/inaccurate content to either experience in a non-arbitrary way, one is not in a position to identify how E1 and E2 differ in content.
- 11. So one is not in a position to identify how E1 and E2 differ in content [9,10].
- 12. So, E1 and E2 differ in phenomenology, yet representationalists cannot specify how E1 and E2 differ in content [2,11].
- So, representationalists cannot explain how the phenomenal difference between E1 and E2 is to be accounted for by a difference in content [12 & Representational Intelligibility].²⁸
- 14. (13) is contrary to assumption (1) [reductio].

As there are other premises to reject to avoid the reductio without giving up representationalism, a brief comment on the premises: Clearly, (7) is a crucial premise and I shall return to it. (5) is a specific commitment that is not obligatory for representationalists, but it seems that the argument can also be run allowing for less determinate contents (cf. Section 4.4 and Stazicker and Nanay on determinacy; see also J. Stazicker 2011).²⁹ I take (8) and (10) to be plausible conditions regarding the assignment of content if one accepts the influences of attention. Finally, there are the empirical premises which can be questioned. Indeed, I have suggested that the idea of attention being something that can be apportioned in the way assumed by (4) is worth critical reconsideration. One might allow that neural modulations will be varied as (4) maintains without holding that attention is thereby varied. In fact, a version of the argument will still work if (4) is replaced with the less loaded claim that, when one shifts attention, there is a variegated modulation of neural activity across the visual field,

along with the assumption that changes in neural activity can change experiential content (e.g., assuming a version of CRP). I shall discuss this version at the end of this section.

Let us focus on premise (7) which expresses

the problem for the idea that there is any distribution or level of attention that entails either veridicality in normal circumstances or illusion. The problem for this view is that there is no way to pick which distribution of attentional resources engenders veridical perception and which engenders illusion.

(Block 2010, 45)

A first response is to ask why there should be any level of attention that "entails" veridicality. Is there any level of illumination that "entails" veridicality, say, for the visual experience of shape? What visual contents are available will depend on attention as well as illumination in complex ways, but no specific level of either will entail veridicality. The point can be dropped without affecting the argument. Problems for representationalists arise once they allow that attention affects content and that attention is distributed in variegated ways across the visual field. Block is certainly right that, from the armchair, any value representationalists choose for n or m will be arbitrary. For assume that some value for attention v is what allows for accurate content p. Given the distribution of attention across the visual field and over time (premise (4)), any shift in attention, which happens constantly, will lead to a landscape of inaccurate perception, namely at locations where the amount of attention is not equal to v. Consequently, much of what one perceives will be inaccurately represented. Moreover, for any v, a small shift from v shifts us from accuracy to inaccuracy, and this seems uncanny. Why, Block asks, should the value v be tied to accuracy in this way?³⁰

One might respond that the issue is empirical. It will take careful experiments to nail down what the value of v is, and this depends in part on identifying the determinants of content. Consider naturalistic theories of content that aim to explain how a mental state has content p in terms of some naturalistic property N. To take a simple case, one might try to explain why a symbol S refers to object O by identifying some appropriate causal or nomological relation between S and O (Fodor 1991), some teleological relation (Millikan 1984), or some mix of the two (Dretske 1995). Obviously, I cannot consider Block's challenge in light of all of these

theories, but I shall raise a challenge for naturalistic accounts of content that endorse CRP.

Let us understand Block's challenge as the claim that the link between neural content and perceptual content is partly a function of the amount of attention deployed across the visual field. The Datta and DeYoe study Block appeals to in premise 4 suggests that this distribution of attention is not uniform, and presumably, given rapid shifts of attention across time, the influence of attention on neural activity will be guite varied. So, not only will perceptual content be a function of distributions of attention, as the Carrasco premise suggests, but neural activity will be as well. Yet neural activity is presumably a determinant of neural representational content. There is, then, a challenge to representationalists who endorse CRP, for CRP is part of a program of accounting for phenomenology ultimately in terms of neural content. This assumes, however, that we can fix neural content so that it can function to explain phenomenology. A version of Block's challenge is that, given facts about attention, we cannot fix neural content in the requisite way, namely, such as to attribute a specific content to a given neural state. But why think what the challenge alleges is true?

Can we identify a neural property that fixes neural content? Specifically, is there a property of a neuron that allows us to determine what it represents? Consider Figure 2.5 which presents the firing rate of a neuron to lines of different orientations in its receptive field and identifies a maximum firing rate for a given orientation, O. Does the neuron in question thereby represent orientation O? Similar questions can be raised for a population of neurons that can be said to represent some feature or object: What neural properties of the population determine what the population represents? Sticking with the simpler single neuron case in Figure 2.5, the neuron's maximum firing rate with and without attention was the same, centered on line orientation O. Let us associate the neuron's representational content with its preferred stimulus as identified by its strongest response, i.e., its maximum firing rate. So, this neuron represents O. Will that response hold at every level of attention, assuming that attention can be graded in this way? What if the maximum neural activity shifts at different amounts of attention? Does what the neuron represents thereby change every time attention shifts? Let us allow this sort of "representational plasticity" in neurons, given shifting attention: what the neuron represents changes with attention. How then does the brain maintain representational stability across time given constantly fluctuating attention? Neurons, it would seem, will "say" different things depending on how attention modulates them. Yet, if a person constantly changed stories due to external pressures, it would be hard to trust the person as an information source. One wants representational stability.

If representational plasticity is unacceptable, then we might respond by tying neural representation to the maximum firing rate at a specific level of attention. But this raises a version of Block's question: which level of attention determines (veridical) content? Adherents to CRP will hold that perceptual content is a function of neural content, but if attention is a free agent here, then it is a good question how we can fix neural content when neural activity is so variable in light of attention. Yet if there isn't any specific level of attention that fixes veridical content—any choice being arbitrary—then the challenge to CRP representationalism is that the anchor for phenomenology is in fact not an anchor at all. To understand perceptual content, the adherents of CRP want to tie it to a clear notion of neural content, but attention seems to muck this up.

It is not clear that the current challenge is insurmountable, and, as it stands, it is only the sketch of a challenge, something that must be fleshed out in further discussion. But the sketch suffices to highlight the fact that theories of content, whether those in philosophy or in cognitive science, have too long operated without consideration of a major influence on the content of a given perceptual or neural state. A central theme of this book is that attention is of rich philosophical significance, and if the current discussion points in the right direction, then attention is of great significance to theories of content. A major project will be to formulate theories of perceptual content with an eye towards the fact that we are creatures of (shifting) attention, and that how we attend can have striking modulations on how things appear to us.

4.8 Summary

Is attention a specific way of being conscious? I began with the Jamesian idea that perceptual attention is essentially a specific way of being conscious. The possibility of unconscious attention led us to restrict the claim to the deployment of attention in perceptual experience, and to search for some uniform feature that serves to individuate conscious attention as a form of consciousness. Our attempt has at best been equivocal. Many empirically supported changes in phenomenology resulting from attention do not seem to fall under one general phenomenal category such as increased determinacy. This leads to a central question: Have theorists

misconstrued the truth in James' original description? Perhaps attention does not have a uniform phenomenology associated with it. Visual attention does not uniformly make things look brighter, of higher contrast, or of greater spatial determinacy. This allows that visual attention (or perceptual attention) can have a variety of effects on conscious experience, but if there is a uniform feature of attention, perhaps it is not that attention has its own distinctive phenomenal upshot, but that, in light of the phenomenology of conscious experience, attention allows for a selective response to it, to focalize and concentrate what is relevant to the subject. Even if there is a phenomenology of attention, what matters in respect of conscious attention is not this phenomenology but the selectivity of attention in respect of consciousness. Attention is conscious to the extent that it is a specific response to consciousness.

Suggested reading

Kentridge (2011) reviews evidence for unconscious attention. A nice overview of the empirical work discussed at the beginning of the chapter, from a neuroscientific perspective, can be found in Anton-Erxleben and Carrasco (2013). Chalmers (2004) provides an overview on representationalism and discusses different aspects of the theory; for a recent monograph discussing these issues, see Siegel (2010). Block (2010) provides a summary of some of Carrasco's work and presents his challenge to representationalism. Stazicker (2011), Wu (2011c) and Watzl (2011) provide recent accounts of attentional phenomenology. Speaks (2010) suggests that attention might have a sui generis phenomenology and presents one of the earlier responses to representationalism in light of attention.

Notes

- 1 The term "what it is like" refers to phenomenal consciousness (Nagel 1974).
- 2 Of course, talk of the essence of attention has modal implications, namely that attention is *necessarily* conscious. How does one establish that claim? Induction can point to some form of nomological, i.e., lawful, necessity, but there is an open question theorists of attention should raise: What precisely are the arguments for James' essence claim?
- 3 That blindsight patients are phenomenally blind in certain regions of their field of view is now well accepted. Skeptics often wonder whether blindsight is

just severely degraded vision (cf. Overgaard 2011 for a recent example, but see earlier work addressing this question using signal detection theory, e.g., Azzopardi and Cowey 1997). I'm grateful to Bob Kentridge for discussion on these and related issues.

- 4 Given that the inputs that inform GY's guesses are likely cortical inputs, albeit abnormally processed ones, relative to what one typically accesses, we should be careful in inferring too much from talk of "guesses". It is not surprising that patients with blindsight should be hesitant to comment on items to which normal access is not available.
- 5 Jesse Prinz has provided a response to the experiments with GY, but I shall assess his response in Chapter 5. To anticipate, I think his objections can be answered.
- 6 The passage is translated by James. Fechner interestingly further notes that "everyone, on the contrary, feels the increase as that of his own conscious activity turned upon the thing" (James, op. cit. p. 426). I shall return to a version of this response in discussing phenomenal salience and one's awareness of one's attending to a specific target.
- 7 That is, for luminance, I, the Michelson contrast, $C_m = (I_{max} I_{min})/(I_{max} + I_{min})$.
- 8 This is reflected in calculations of the *point of subjective equality* (PSE), namely the point at which the test and standard appeared to be the same contrast. I shall not be concerned with this technical detail, and I simply report the result of such calculations in each experiment.
- 9 It is worth noting that Massimo Turatto et al. (2007) found that attended *moving* Gabor patches appeared of *lower* contrast.
- 10 For an alternative explanation based on a shift in decisional criteria, see (Schneider and Komlos 2011). For responses, see (Anton-Erxleben, Abrams, and Carrasco 2010) and (Anton-Erxleben, Abrams, and Carrasco 2011). For an alternative explanation based on cue bias, see (Prinzmetal, Long, and Leonhardt 2008) and for a response, see (Carrasco, Fuller, and Ling 2008). Note that if the criticisms of Carrasco's account are correct, this does not hurt the case against the phenomenal conception. We are trying to find a way of characterizing attentional phenomenology grounded in empirical work. If the current avenue is closed, so much the worse for the phenomenal conception. One can, however, raise a question for Carrasco, since her experiments draw on the very short presentation times of the central stimuli. Yet another well-known paradigm due to George Sperling explicitly explores what we can be conscious of in short presentations of stimuli (roughly on the same time scale, if not longer, than

that of the presentation of stimuli in Carrasco's experiments). In that case, many questions have been raised about what can be seen (Chapter 6).

- 11 Later work also showed changes in the appearance of spatial frequency, namely increasing apparent spatial frequency, during bottom-up attention (Gobell and Carrasco 2005) and top-down attention (Abrams, Barbot, and Carrasco 2010).
- 12 This is derived from "Landolt C" figures that are used in European countries in visual acuity tests (a "C" is placed in different orientations, with subjects asked to identify the location of the gap, e.g., up or left).
- 13 Carrasco and colleagues argue that the best explanation for these results is that attention enhances spatial resolution. For a brief summary of alternative explanations, see for example (Montagna, Pestilli, and Carrasco 2009, sec. 4.5). For more in depth discussion, see Carrasco 2011. How attention enhances spatial resolution so as to explain this task is still an open question. Chapter 2 noted that attention can lead to contraction of visual receptive fields around the target of attention, and smaller receptive fields are associated with better resolution. Still, moving from this to the perception of gap size is less clear. Katharina Anton-Erxleben, Christian Heinrich, and Stefan Treue (2007) have observed that attention increases the apparent size of moving patterns of dots, and they make suggestions of how remapping of receptive fields might lead to distortions in perceived visual space. See also Anton-Erxleben and Carrasco 2013 for recent discussion.
- 14 The changes in experience that are induced by voluntary shifts of attention are reminiscent of the changes that can be induced in the experience of ambiguous figures such as the Necker cube or the duck-rabbit figure (Tse notes this as well). That is, shifts of attention seem to allow one to flip the orientation of the Necker cube or whether one sees the figure as a duck or a rabbit.
- 15 In the text, I emphasize that attention causes visual phenomenology. In this way, the relevant phenomenology is too distant from attention. One might, however go adverbial and speak of attentive and inattentive visual experience, i.e., a way of visually experiencing the world (recall Mole's adverbial account). If so, then the relevant phenomenology's being a property of experience need not be too far removed from attention. If the phenomenology can be disconnected from attention, however, this move will fail. See the next section for such an argument.
- 16 Changes in determinacy might be general enough to cover the effects of attention across all sensory modalities. Stazicker and Nanay differ, in print at least, on how extensive they take the Carrasco effects to be. Stazicker is

agnostic as to whether all attentional phenomenology involves increases in determinacy. Rather, he thinks an important aspect of the phenomenology of visual attention is its mimicing the effects of foveation, as when you shift your eyes to look directly at something. Foveation improves spatial resolution, and attention often does as well.

- 17 See also Block 2010, 43.
- 18 Hemdat Lerman (manuscript) has raised similar questions, suggesting that when one reflects on one's attention while tracking multiple objects, there isn't any clear attentional phenomenology that accompanies each attended object (for multiple object tracking, see Chapter 7). She notes that much of the current theorizing about the phenomenology of attention has focused on a limited set of static samples that are quite restricted, as when one reflects on Carrasco's Gabor patches and thinks about how attention affects experience.
- 19 Watzl suggests that one can further explicate these ideas, if we think of experience as composed of parts. Thus, an experience of Figure 4.2 contains as parts the experience of the 22% contrast Gabor patch (E_G), as well as the experience of the fixation point to its right (E_F). Watzl expresses the structure of experience as follows when one attends to the Gabor patch G: $E_G > E_F$. That is, E_G is more central than E_F . The idea of structuring allows for multiple centers that each is the focus of attention. Attention then could be divided in the sense of having multiple foci of attention.
- 20 There is much more in Watzl's view that we cannot discuss here. See his forthcoming book *Attention and the Structures of Consciousness*.
- 21 There are complications here, depending on whether one endorses *pure* or *impure* representationalism (Chalmers 2004). Pure representationalists appeal only to the content of the conscious state; impure representationalists also appeal to the psychological mode of the state. I shall consider pure representationalists here. Some representationalists endorse a different supervenience claim, namely, that representational content supervenes on phenomenal character but not *vice versa*. Specifically, I shall focus on many of the earlier forms of representationalism that endorse the supervenience claim in the text.
- 22 For an earlier discussion of different versions of representationalism, see Byrne 2001.
- 23 For an overview of different positions on perceptual consciousness, see the early sections of Pautz 2010.
- 24 Fiona Macpherson presents a different case against representationalism by appeal to ambiguous figures like the duck-rabbit. Although she does

not discuss attention explicitly, it seems that attention can play a role in these *gestalt shifts*. As the phenomenal changes at issue seem to primarily concern gestalt shifts that can be induced by attention, rather than as what one might be inclined to call attentional phenomenology, I have opted to not discuss the issue at length here. I am grateful to Guilia Martina for some helpful pointers on this literature.

- 25 Jeff Speaks (2010) presents another case to which Carrasco's results might readily apply.
- 26 Block's assumption that Carrasco's observation can be localized to a single Gabor patch can be questioned. Block is aware of the need to justify this assumption, and in support, he notes Carrasco, Penpeci-Talgar, and Eckstein 2000, which focuses on boost in the visual signal for individual Gabor patches. It is worth noting that this work uses a different paradigm than the one Carrasco uses to test appearances in Carrasco, Ling, and Read 2004; in Carrasco's experiments, the task involved judging relative contrast (Which is higher in contrast?). Thus, any relevant illusion concerns the appearance of relative contrast, and not the specific contrast of a particular Gabor patch. Nevertheless, one might think that once we allow that attention can alter contrast appearances to affect the appearances of relative contrast? The issue here is that the argument draws on this point given empirical work, but it is plausible that this work only speaks to relative, and not absolute, contrast.
- 27 Block also emphasizes that this variability is not only spatial, but also temporal, though we will not emphasize that. Adding the temporal element compounds the problem of shifts in content with shifts in attention, now over time as well as space. If we think of attention as a limited resource, then any changes over time in the amount of attention at one region necessitates changes in the amount of attention at another region.
- 28 Sebastian Watzl (forthcoming) presents a detailed analysis of Block's argument. Watzl holds that both E1 and E2 are illusory. He emphasizes that we should construe the function of attention as providing for *usable* and not necessarily *accurate* representation, unless accuracy serves usability. Boosting a signal might be the way that attention aids the usability of perception, but perhaps at the cost of accuracy. A selection for action account would find this response congenial. The function of attention is action, not veridicality. At the same time, since epistemic actions (e.g., fixing beliefs) are important, too, attention could not wildly introduce inaccuracy on pain of undercutting one's cognitive goals.

- 29 In a personal communication, Ned Block emphasizes that it is this premise that he ultimately rejects. He allows that there are accurate contents, but that these contents do not account for the phenomenology. Two brief points: (1) it seems to me that Block's argument questions whether we can assign a specific content to experience, whether that content is determinate or not. Given the effects of attention, how would we fix a content?; (2) Block claims that if you allow for indeterminate content, phenomenology doesn't flow from it. Things don't look indeterminate. I say: they do. Say that you have normal vision. If you were to go to the optometrist office and have your eyes tested, the optometrist can use lenses to slightly blur your vision. So, presumably your experience has less determinate content. But was it fully determinate before hand? I think not. But then phenomenology that might seem determinate is in fact indeterminate. That is, how things look to you now is precisely how things look when your content is indeterminate. No one has fully determinate spatial vision. What would that be like? Fully determinate phenomenology is what might be mysterious. I'm grateful to Jake Beck for discussion here (the eyeglass example is his). For some relevant discussion, see his (2012).
- 30 We can escape the pervasive inaccuracy in perception due to shifts of attention if we allow that content is less determinate than premise (5) asserts.

5

ATTENTION AS THE GATEKEEPER FOR CONSCIOUSNESS: INATTENTIONAL BLINDNESS

5.1 Introduction

William James (2007) comments as follows on attention and conscious experience:

Millions of items of the outward order are present to my senses which never properly enter into my experience. Why? Because they have no *interest* for me. *My experience is what I agree to attend to*. Only those items which I *notice* shape my mind – without selective interest, experience is an utter chaos. Interest alone gives accent and emphasis, light and shade, background and foreground – intelligible perspective, in a word. It varies in every creature, but without it the consciousness of every creature would be a gray chaotic indiscriminateness, impossible for us even to conceive.

(402-3, James' italicization)

"My experience is what I agree to attend to" suggests that attention serves as a *gatekeeper* for consciousness. It determines what one is conscious of. This chapter begins an extended discussion of attention as gatekeeper. To see the distinctiveness of the gatekeeper model, one can contrast it with a more common sense model where attention selects from the deliverances of consciousness (see Figure 5.1).



Figure 5.1 The Gatekeeper and Common Sense Models

I shall initially characterize the gatekeeping model as endorsing the following biconditional:

Gatekeeping: S is conscious of X if and only if S attends to X.¹

This model provides a different answer to the metaphysical question: attention is a selective process that is for consciousness. In other words, attention is identified by its distinctive functional link to consciousness. What one attends to determines what one is conscious of. Gatekeeping has been a subject of significant scientific work in the past two decades, and this specific issue kept attention alive as a topic for philosophical discussion during that period. This chapter shall focus on the role of perceptual attention as gatekeeper, clarify the central claims, survey the empirical evidence for gatekeeping, and assess the state of play. The central questions are:

- What does it mean to say that attention is a gatekeeper for consciousness?
- Is there evidence for or against attention as gatekeeper?

Section 5.2 identifies numerous gatekeeping conditionals. That there are so many gatekeeping claims underscores the need for theorists to be clearer about the position they defend. Failure to do so has tripped up some discussions of gatekeeping (see Section 5.7.1). In Section 5.3, the sufficiency claim of gatekeeping is considered and rejected. Section 5.4 then summarizes the standard experiments in favor of the necessity claim: one is conscious of X only if one attends to X. This work leads to the claim of inattentional blindness in the visual domain. Section 5.5 explores what might be meant by "blindness," and three notions of blindness are examined. Section 5.6 then explores what follows from inattention, identifying three alternatives to inattentional blindness: inattentional amnesia, inattentional

agnosia, and inattentional apraxia. In particular, inattentional apraxia suggests that the experimental paradigms standardly used to support inattentional blindness fail to provide evidence for it. Section 5.7 explores attempts to demonstrate consciousness outside of attention, e.g., in the perception of the gist of a scene and in the visual phenomenon of crowding.

5.2 Gatekeeping conditionals

Attention has often been thought of as for consciousness. On Anne Treisman's Feature Integration Theory (Chapter 1.6), for example, attention serves in vision to bind feature representations to generate object representations and, as Treisman notes, *awareness* of objects. In this chapter, I shall focus on perceptual attention and perceptual consciousness, though I will generally omit the "perceptual" modifier and speak simply of attention and consciousness. As noted in the previous section, the gatekeeping conception of attention can be initially unpacked in terms of two claims that together form the Gatekeeping biconditional:

- 1. If S is attending to X, S is conscious of X.
- 2. If S is conscious of X, S is attending to X.

Attention to X, in other words, is both sufficient and necessary for consciousness of X. For current purposes, "conscious of" and "attention to" are success terms. Both imply the existence of the relevant target. Now (1)might seem to be true, for whenever you attend to an object, you seem to be conscious of it. In contrast, (2) might seem to be false: can't one be conscious of more than one attends to? Interestingly, there is empirical work that suggests that (1) is false, but (2) is true.

(1) and (2) provide the most natural interpretation of gatekeeping, but two complications arise once one acknowledges the many distinctions regarding attention discussed in earlier chapters: first, the different targets of attention and consciousness; and, second, the dichotomies of attention such as top-down versus bottom-up and control versus automatic attention. If one incorporates these distinctions into a characterization of gatekeeping conditionals, one generates a slew of additional gatekeeping claims. First, notice that the variables in (1) and (2) take the same argument, so attention and consciousness have the same target. Thus, they suggest that to be conscious of a bear, you have to attend to the bear, and when you attend to the bear, you are conscious of it; to be conscious of the blackness of its fur, you have to attend to the color and where you attend to the blackness, you are conscious of it. But attention and consciousness can have different targets, and to cast the conceptual net as widely as possible, one should allow for different arguments for the variables, thus transforming (1) and (2) into

3. If S is attending to X, then S is conscious of Y, where possibly $Y \neq X$. 4. If S is conscious of X, then S is attending to Y, where possibly $Y \neq X$.

For example: to be conscious of a bear, you attend to some feature of it or where you are attending to some feature of a bear, you are also conscious of the bear. Consider, by way of example, the different gatekeeping conditionals in respect of (4) for consciousness of objects and attention to different targets:

A. If S is conscious of object O, then S attends to O.

B. If S is conscious of object O, then S attends to (some) feature F of O.

C. If S is conscious of object O, then S attends to location L of O.

Does consciousness of an object imply that one is attending to it, or at least to one of its properties or its spatial location? If one allows the variables in (3) and (4) full freedom to range over location, feature or object, then for each of (3) and (4), nine distinct gatekeeping conditionals result. This is dizzying. The diversity of conditionals raises an important question: which conditional(s) are at issue in discussions of gatekeeping? As an exercise, look at any of the articles cited in this chapter and see if the authors explain clearly which conditional they aim to defend or refute.

Briefly, let us note the second complication, namely adding the distinction between top-down versus bottom-up and control versus automatic attention into the fray. Consider the following three conditionals:

- 5. If S is conscious of X, S is top-down attending to X.
- 6. If S is conscious of X, S is bottom-up attending to X.
- 7. If S is conscious of X, S is top-down or bottom-up attending to X.

(7) is the most general since the dichotomy of top-down versus bottom-up as defined in Chapter 1 is exhaustive. So (7) is just a version of (2): consciousness of an object requires some form of attention to it. (5) and (6) are more restricted theses, but I bring this up to highlight their distinctness.

Let us, however, leave these complications aside. In noting them, it is clear that one must treat talk of gatekeeping with care. Having acknowledged this, I return to (1) and (2).

5.3 Attention is not sufficient for consciousness

Chapter 4 presented experiments that suggest that (1) is false. Recall that using the Posner spatial cueing paradigm, Kentridge et al. (Kentridge, Heywood, and Weiskrantz 1999) demonstrated that the blindsight patient GY showed standard attention cueing benefits to stimuli in his blind field, suggesting that GY could attend to objects that he could not consciously perceive.² In those experiments, a cue directed GY's attention into his blind field, and he reported on the presence or properties of the target therein. GY showed improvements in reaction time and accuracy for valid versus invalid cues that in normal subjects are a sign that attention is engaged. Since attention is being deployed in the blind field, the claim is that GY demonstrates attention to X without consciousness of X.

Jesse Prinz (2012), who endorses both (1) and (2), has objected to this interpretation of the result, so let us consider his objections.³ He notes that the cueing paradigm is a test of spatial attention, but the relevant form of consciousness is of an object. Thus, the relevant sufficiency claim is:

If S attends to object O, then S is conscious of O.

Prinz notes that the cueing paradigm only establishes attention to a location, not to an object. Thus, the experiment does not give a case where the antecedent is satisfied, so it does not present a counterexample to (1). In this way, he is being appropriately sensitive to different gatekeeping conditionals. Nevertheless, spatial cueing involves object attention as well. After all, to perform the task, the subject must issue a report about a specific object, and not about a location. It is a strange position to allow that attention gets pulled by the cue but then simply leaves the scene when the target to be reported on appears. Rather, attention is also needed to select the target to inform the subject's report as the task requires. In many cases, such selection for task suffices for attention. In reporting on the presence of a target, GY thus attends to an object that he cannot consciously see. Prinz is correct that spatial cueing is deployed as a test for spatial attention, but the standard construal of the paradigm misses the essential involvement of object attention.

Prinz's central response is to distinguish orienting from attention and to argue that the putative counterexamples to (1) are due to orienting. But what is orienting? Prinz writes that "orienting alters what information gets in and attention alters where it flows" (p. 113). In vision, overt orienting aims to allow a target of interest to stimulate the fovea, the retinal area of highest visual acuity. Thus, Prinz ties overt orienting to what is typically called "overt attention," attention linked to movement of the sensory organ. He also allows for two forms of *covert* orienting: (a) when the subject forms an intention to overtly orient; or (b) when there is "shrinking of receptive fields" (p. 114).

Can invoking covert or overt orienting provide an adequate defense in response to the experiments with GY? I do not think overt orienting can. Prinz suggests that GY might be making microsoccodes, smaller saccadic (ballistic) movements of the eye. Since saccades result in foveation, namely allowing a target to stimulate the fovea, Prinz's idea must be that microsaccades provide enhancement of visual processing in the direction of foveation. Basically, GY is sneaking a peek. While I agree this is possible, I think physiological parameters speak against it. While the upper bound of microsaccades is often taken to be one or perhaps even two visual degrees, microsaccades are typically much smaller, say, less than 30 arcminutes (1/2 of a visual degree; Rolfs 2009). In GY's case, however, the target was located six visual degrees from fixation, and experimenters could detect if GY made eyemovements of two to three degrees. I presume that they threw out any trials in which such visible movements were made. Prinz would then need to explain how microsaccades can explain enhanced performance of a target likely located three, and more likely at least six, times the distance of a standard microsaccade movement from fixation. Microsaccades do not explain GY's performance.

What of covert orienting? For Prinz, one implication of covert orienting is visual field remapping, i.e. the phenomenon of receptive field remapping discussed in Chapter 2: when two objects are present in a neuron's receptive field, attention to one of them results in remapping the receptive field to the attended object, almost as if the receptive field shrinks around it, filtering out the other object. As discussed in the next chapter, Prinz proposes that attention is selection for working memory, but his proposal is relevant here. The experimental task used to induce the effect in the first published study by Moran and Desimone (1985) required the use of working memory. Animals were tasked with matching a stimulus with an earlier sample. Moran and Desimone suggest that the observed effect is a kind of neuronal filtering of irrelevant stimuli that supports working memory. By Prinz's own characterization of attention, receptive field remapping is tied to attention as such remapping serves working memory. Thus, while Prinz invokes covert orienting as distinct from attention, covert orienting as receptive field remapping implicates attention.

Finally, what of covert orienting as intention, i.e., an intention or at least preparation to move the eye (saccade)? This is the most promising defense. Prinz's argument would be that: (A) covert orienting as an intention to move the eye is not the same as attention; and that (B) the intention modifies processing sufficient to enhance GY's performance. Neuroscientists have long debated the precise relation between orienting qua intention to move and attention. For example, at the beginning of electrophysiological recordings with awake behaving monkeys, there was a heated debate whether one can even distinguish intention from attention in neural activity (see Glimcher 2004, chap. 10). Recall as well the Premotor Theory of Attention that proposes a constitutive link between the preparation of eye movement and attention (Rizzolatti, Riggio, and Sheliga 1994; see also Chapter 2).

As noted in Chapter 2, there are compelling reasons to reject strong versions of the Premotor Theory and thus to endorse (A), but it seems that many scientists would allow that preparation of an eye movement (or an intention to so move) towards X often causes attention-based enhancements of processing at X (see Smith and Schenk 2012, sec. 3; Armstrong 2011). Indeed, it is arguable that any intention to move the eye to location X requires prior attention to X to make possible an intention with that very content. In other words, covert attention is needed to provide the targeting needed for orienting. The point, then, is that the relevant intention to move the eye might both cause and depend on attention. While I think Prinz's appeal to intention is interesting, the fact that intentions are closely tied to attention suggests that attention is at least as likely, and to my mind more likely, an explanation of GY's performance as intention. Prinz's objections can be met. So, given the empirical evidence from GY and from normal subjects (see Chapter 4.2), (1) seems to be false. Attention does not suffice for consciousness. If (1) is false, then attention as selection for consciousness does not provide a single, complete answer to the metaphysical question, since not all forms of attention are for consciousness. Accordingly, the gatekeeping thesis is restricted to (2): one is conscious only of what one attends to. Is there evidence for or against (2)? Surprisingly, the question remains quite open.⁴

5.4 Attention as gatekeeper for consciousness: initial evidence

To assess the gatekeeping view, consider conditional (2) now restated as the Gatekeeping Hypothesis:

Hypothesis_(GK): S is conscious of X only if S attends to X.

We can state more specific versions of Hypothesis_(GK) for relevant sensory modalities. Thus, S is visually conscious of X only if S visually attends to X. What happens when you fail to attend to X? In the case of vision, you will not be phenomenally conscious of X even if it is in your field of view. In this case, cognitive scientists speak of inattentional blindness and change blindness, terms that derive from a set of experimental paradigms. Let us begin with cases of inattentional blindness. If you have never seen displays demonstrating this, I don't want to spoil it for you. Stop reading and follow the link to the "Monkey Business Illusion" (at http://www.dansimons.com/videos.html). Alternatively, do an Internet search for "Monkey Business Illusion." Be sure to follow the instructions. Even if you are familiar with the paradigm, try it again with this version. Do it now and don't read the next paragraph until you do!

The video demonstrates a form of inattentional blindness: by directing your attention to the ball and the number of passes, your attention is pulled away from the gorilla. Moreover, you don't seem to be conscious of the gorilla in that you don't report seeing it. When shown the video again, you are surprised that you failed to see it.

Arien Mack and Irving Rock (1998), who coined the phrase inattentional blindness, used a different paradigm. A cross is flashed for 200 milliseconds (ms), followed by a mask to disrupt residual visual processing of the cross (e.g., an afterimage). Subjects were required either: (a) to report which arm of the cross was longer. or, in some trials, (b) whether they were equal in length. After a number of trials where the cross is flashed at the same location in the visual field, an additional stimulus was presented with the cross in the critical trial. Subjects were never apprised of the possibility of an additional stimulus. After the critical trial, subjects were then asked if they saw anything in addition to the cross, something they had not seen in previous trials. Strikingly, many denied seeing anything new. Mack and Rock (1998) summarized their findings as follows: All the experiments described in [our] book demonstrate that when one takes the proper measures to eliminate voluntary attention or, better expressed, the intention to attend to what is about to be displayed, there is a drastic reduction in what is perceived, at least on a conscious level. In fact, unless certain kinds of objects are presented in our critical trial, the rule seems to be that *nothing is perceived consciously*.

(163, my emphasis)

In other words, Mack and Rock endorse attention as a necessary condition for consciousness, the gatekeeper view. Here, the focus seems to be on topdown, goal-directed attention, but we shall see later in Section 5.7 that they also have bottom-up attention in mind.

Mack and Rock found that the level of inattentional blindness, as measured by the percentage of subjects who failed to report the additional stimulus in critical trials, varies with the type of stimulus and with its location within the visual field. When a solid black geometrical shape was flashed in the periphery while the cross was flashed at fixation, 25% of subjects denied seeing the black shape. Strikingly, when the cross was flashed in the periphery and the black shape placed at fixation, 60 to 80% of subjects denied seeing the shape! This is surprising given that the unexpected shape stimulates the fovea. Finally, when the stimulus was the subject's own name, inattentional blindness dropped to 12.5% of subjects (Mack and Rock 1998, fig. 5.1, p. 117). While I will later argue that these experiments cannot provide evidence for inattentional blindness, they do provide a wealth of information about what sorts of stimuli capture attention and thereby interrupt ongoing task performance (recall the idea of bottom-up attention as a circuit breaker, Chapter 1, Section 7).

Change blindness is different in that it is defined in terms of a specific target, an *event*. Standard change blindness paradigms involve flashing two nearly identical pictures in succession with an intervening mask (a gray screen), cycling from one to the other. As the second picture is different in some way from the first, there is a change between the two. Subjects are asked to report the change. What is striking is that despite knowing that something is different, subjects often take quite a bit of time to spot the change, even if the change is substantial involving large portions of the scene (e.g., a whole building disappearing). Inattentional blindness is the more general phenomenon, with change blindness as a species.

Note that there are, at least, two other phenomena that have led to talk of inattentional blindness: hemispatial neglect and the attentional blink. Hemispatial neglect involves a patient's unawareness of a certain side of space, typically the left side, due to damage to right parietal cortex. Strikingly, neglect patients might fail to dress the left side of their body, fail to eat from the left side of their plate, or fail to report on the left side of objects. Neglect is often described as an attention deficit, an inability to attend to the left (or relevant) side of space (Losier and Klein 2001; Bartolomeo, Schotten, and Doricchi 2007). The attentional blink can be demonstrated in normal subjects and is characterized by the inability to notice a second stimulus rapidly presented after an initial stimulus, as if after attending to the first, attention then blinks and misses subsequent stimuli within a small temporal window (Martens and Wyble 2010; Dux and Marois 2009). In both cases, it is postulated that the inability to attend to a stimulus implies that the subject is not conscious of the stimulus.⁵ We do not have the space to discuss these interesting phenomena, but the criticisms I shall later make of inattentional blindness experiments can be applied to these cases as well.

This empirical work suggested to many cognitive scientists that the absence of attention to an object is sufficient for the absence of consciousness of that object. In the visual domain, inattention suffices for blindness, i.e., inattentional blindness. Perhaps this seems to you to be the right conclusion, but there are two major questions that must be answered before one can understand what the claim is: What is meant by "blindness," and what conception of attention is in play?

5.5 On blindness

What is meant by "blindness" in Hypothesis_(GK)? It can't just mean the mundane idea of a failure to notice, as when one yells, "Are you blind?", at a friend who fails to notice something visually obvious. The inattentional blindness paradigms demonstrate blindness in that sense, but such blindness is compatible with actually visually experiencing what one has failed to notice. That is why we think our friend is subject to criticism for his inattention. Such blindness is unfortunately all too common. Accordingly, for the thesis of inattentional blindness to be interesting in respect of phenomenal consciousness, "blindness" must be taken literally. But what does this come to? Surprisingly, there is almost no discussion of this in a literature that constantly speaks about blindness. Theorists generally speak of the lack of consciousness of an X, but what does this really come to? This is a glaring lacuna: theorists make provocative claims about the relation between attention and consciousness, yet the charge of blindness is never

adequately explained. I have already raised the mundane notion of blindness as failure to notice, so if a stronger sense of "blindness" is intended, what is it? As a way to isolate different answers, let's keep track of the following gatekeeping conditionals formulated as sufficient conditions for blindness:

- **Object blindness**: failure to visually attend to an object O at location L suffices for blindness to O.
- **Feature blindness**: failure to visually attend to an object O at location L suffices for blindness to O's features.
- Location blindness: failure to visually attend to an object O at location L suffices for blindness to L.

The congenitally blind are, of course, blind in all three senses. The question then is this: When a theorist claims that subjects are inattentionally blind, which of these conditionals do they intend when they speak of blindness?

Invocation of the congenitally blind suggests a natural interpretation of "inattentional blindness," namely that inattention leads to phenomenal blindness or the absence of any visual experience: one is blind to features, objects, and space.⁶ Blindness then identifies the absence of sight. Phenomenal blindness is not darkness, as when you close your eyes, for experiencing darkness is a visual experience not available to the congenitally blind. One tells by vision whether it is dark. To get a sense of what it is like to be literally blind, fixate on some object before you, and then hold out your thumb and slowly move it to the periphery of your field of view while you maintain fixation. At some point, you will no longer see the thumb. At the point at which the thumb disappears, it does not disappear into darkness. Rather, it just disappears into an area that you cannot from the current vantage point visually experience. Having a normal field of view is not like looking through a tube, an area of visibility with a penumbra of darkness. When your thumb disappears, you are literally blind to it: you have no visual experience of it. At the limit, where the visual field disappears, you do not have visual experiences at all. To imagine what total blindness is like, imagine the boundaries of the visual field collapsing but without being replaced by darkness. That would be literal blindness. On this reading, the claim of blindness is stronger than the claim of failure to notice. If failure of attention is sufficient for blindness, then a strong version of the gatekeeper thesis takes attention to define the extent of the visual field. Anything to which one does not attend lies outside the visual field, and one has no visual experience of it, whether of space, feature, or property. Thus, inattentional

phenomenal blindness entails the truth of object, feature, and location blindness.

Let us apply this notion of phenomenal blindness to inattentional blindness of the gorilla. The idea would be that since subjects are not attending to the gorilla, they are phenomenally blind to it. Let us assume that there is some configuration where the visual field covers an area including all the white and black clad players passing the balls, and within which the gorilla appears. Given inattention, subjects are phenomenally blind to the object that is the gorilla, its spatial location, and its visible properties. That part of the world lies outside the visual field. How, then, should we characterize the relevant blindness? When the gorilla walks into the scene of the players passing the ball, what happens to the visual field? Initially, a subject observing the players will have a visual field that encompasses the players passing the relevant ball. When the gorilla walks into the scene, perhaps it acts like a magnet, drawing the visual nothingness that characterizes what is outside the visual field with it, as if the boundaries of the field are drawn in with it, constricting the area of the world that is visible to the subject. Now as the gorilla walks among the players into the center of the scene, does the boundary get pulled in further, leaving a trail of blindness behind it? Or does the visual nothingness pinch off like a bubble, so that, as the gorilla walks across the visual field, we have a phenomenal hole corresponding to its location, one that moves across the area where the players are passing the balls. Inattention leaves a hole that is not visible in any sense, just as the areas outside of a normal visual field are not visible in any sense. Indeed, the location of the hole will shift as attention shifts with the ball. But is experience of the world really like that? Is that what proponents of inattentional blindness are committed to? The phenomenal hole version of inattentional blindness is quite radical.

Let us consider a second interpretation of "blindness," as when one speaks of the color blind. Consider the Ishihara color vision test that involves a set of color dots in a circle where some of the dots are differently colored and form a numeral, with the other dots forming the background (see Figure 7.6 for a black and white version of this). For certain colorblind individuals, the numeral cannot be discriminated from its background, and they are in that sense blind to the numeral. They cannot see the numeral because they cannot distinguish the dots that constitute it from the other dots that form the background. Of course, these subjects are neither blind to the dots, which they can see quite well, nor are they blind to color, for the dots can appear colored. The problem is that they do not appear differently colored. Think also of an animal that is camouflaged, e.g., a chameleon that blends in with the bark of a tree. One might see the color of the chameleon without individuating it, contrasting it against its background. In such cases, objects are not perceptually individuated even though some of their features are perceivable.

There is then a conception of blindness in the sense of an inability to differentiate an object from the background. Call this individuation blindness. Accordingly, in inattentional blindness paradigms, one might be blind to the gorilla in the sense that one cannot distinguish it from its background or because, in some sense, only its features are revealed, though the gorilla itself is not (just like the camouflaged chameleon). Let the gorilla then be at location L and surrounded by background B. Blindness to the gorilla is just that the gorilla melds with its background or is not sufficiently resolved. In this sense, one is blind to the gorilla, but notice that some consciousness is preserved in individuation blindness: one remains aware of the location of the gorilla and, perhaps, of some of the gorilla's features, e.g., the blackness of its coat. On this conception of blindness, attention does not define the extent of the visual field. Inattentional blindness in this second sense entails object blindness, but not location or feature blindness. One could be conscious of location and feature even if one is not attending to these, but elsewhere, e.g., to the ball. Similarly, in the Ishihara figure, one might be blind to the numeral, even though one can see the location and certain features of the dots that constitute the numeral. If proponents of inattentional blindness mean individuation blindness, then the gatekeeping conditional is restricted to objects: one is conscious of an object only if one attends to it. Otherwise, inattentional individuation blindness would collapse into inattentional phenomenal blindness.

Finally consider *category* blindness (I will discuss this later as inattentional associative agnosia) where one is blind to the gorilla's being a gorilla. That is, one does not see it as a gorilla. Notice, however, that while visually missing this higher-level categorization suggests a type of blindness to the gorilla, it is compatible with other ways of seeing the gorilla. For example, one might still see the gorilla as a moving black shape, one that task instructions explicitly require one to ignore (follow just the white-shirted, not black-shirted, players). On this version of blindness, there is a selective visual deficit in categorization of an object that is compatible with consciousness of the object in another, lower-level sense. On this version of inattentional blindness, there is *category* blindness but not object, feature, or location blindness.

I have just canvassed three concrete meanings of "blindness" which can have different implications for the limits in visual consciousness under conditions of inattention. Many empirical theorists of this phenomenon seem to have the strongest version in mind, inattentional phenomenal blindness, but it might seem more plausible that there is either individuation or category blindness. The central point is that the failure to be clear about how to characterize blindness has left the basic claim of inattentional blindness unclear.

5.6 What actually follows from inattention?

Does the experimental evidence demonstrate blindness given inattention? The crucial step in all the relevant experiments begins with subjects' failure to report X or denial that they were conscious of X, and then draws on this as evidence that subjects fail to consciously experience X. Since many of the tasks are intended to engage and direct the subject's attention away from X, experimenters further conclude that the subject is not attending to X. Thus, there is evidence both for lack of consciousness of X and lack attention to X. This is followed by an inference to the best explanation: the subjects are not conscious of X because they are not attending to X. This supports the gatekeeping view: failure of attention to X suffices for failure of consciousness to X. The inference, however, requires that inattentional blindness be the best explanation among alternatives. But is it? I briefly identify three alternatives to inattentional apraxia.

5.6.1 Inattentional amnesia

The first alternative hypothesis is that the observed failure to report the critical stimulus is not due to failure of consciousness, but to failure of visual processing that puts the relevant item into memory. Since report of a previously seen stimulus requires memory of that stimulus, failure of storing the relevant object in memory will suffice for failure to report, even if the stimulus was consciously experienced. In fact, this seems to be a natural explanation of change blindness, a form of inattentional blindness. If one thinks of change as a modification in the property of an object over time, say, its transition from exemplifying F to exemplifying G, then it is natural to think that to experience the change is to be aware of the difference between F and G. For example, when one notices that an object is moving,

one is aware of the difference in its spatial location over time; when one notices its changing color, one is aware of the difference between the old and new color. One in some sense compares and contrasts them. Indeed, the use of "comparator" mechanisms that involve comparison between previous and new information has been postulated to explain various phenomena such as efficient motor control (Wolpert and Ghahramani 2000), or why the world appears to be stable when one moves the eyes (Wurtz 2008). On this view, to see the change, one has to remember the old, so as to contrast it with the new. One way a subject can fail to detect the change is to not retain previous information. In that sense, the subject does not store it in memory. The proposal, then, is that in conditions of inattention, one will be more prone to memory lapses. Indeed, notice that in the experiments discussed, the required reports of unattended stimuli are generally retrospective: after the presentation of the critical stimulus, subjects are asked if they noticed anything unusual. If subjects are to accurately report the stimulus, they must remember that the stimulus was present. Since report requires memory, failure of memory can preempt the absence of conscious awareness as an explanation of the subject's failing to report the change. Inattentional amnesia can trump inattentional blindness.

Change blindness is actually both common and familiar. Consider a mundane example. Imagine a case where you meet a man for the first time where he is sporting a mustache. Looking straight at him, there is no question that you are visually aware of his mustache even if you don't make much of it. Five days later, you might run into him again where he is clean-shaven. In such circumstances, you might not be aware of the change and in that sense be change blind. This is not surprising as we often have such experiences even with old friends who might be sporting new glasses or a new haircut. Memory of the past fails us in the sense that we do not have such memories to be compared with the present so as to allow us to notice a change.

Jeremy Wolfe (1999), in responding to Mack and Rock's paradigm, proposed such inattentional amnesia to account for the subject's failure to report the stimuli. He comments:

If vision has no memory and if attention is the gateway to other mental representations, it follows that unattended visual stimuli may be seen, but will be instantly forgotten; hence, *inattentional amnesia*.

(75)

Other researchers in this area have also come to similar conclusions (Lamme 2004; Simons and Ambinder 2005). The argument for the gatekeeping

view assumes that the best explanation of a subject's failure to report X is the subject's failure to be consciousness of X. Yet inattentional amnesia provides an alternative explanation of the subject's failure to report, one that is no worse than the appeal to failure of consciousness. This alternative, if cogent, halts the inference to the best explanation to inattentional blindness.

There is, however, a problem with the appeal to inattentional amnesia as a general response to inattentional blindness. A central assumption in that response is a link between attention and memory, namely, that failure of attention leads to failure of memory. Specifically, the relevant attention at issue is selection for working memory, the sort of memory that is deployed to guide behavior such as recalling a change. The problem is that many theorists postulate that working memory is necessary for consciousness. On their view, inattentional amnesia would be sufficient for inattentional blindness. I will examine the merits of this tight link between memory and consciousness in the next chapter, but in the current context, the appeal to inattentional amnesia is highly controversial. It thus does not provide a clear alternative to the appeal to inattentional blindness.

5.6.2 Inattentional agnosia

A second alternative to inattentional blindness is inattentional agnosia (see Daniel Simons 2000, who attributes the idea to Jeremy Wolfe). Agnosia concerns an inability to see objects or to see objects as belonging to specific categories (see the detailed discussion in Farah 2004). Damage to certain visual areas can lead to deficits such as the inability to see shape and, hence, to see objects (what is often called apperceptive agnosia; Farah 2004, chap. 2). Imagine something like a scrambled visual world (the most well known of apperceptive agnosics is perhaps the patient DF; see Goodale and Milner 2004). Alternatively, some damage to the visual system leads to the inability to see objects as the type of objects they are (categorization) despite largely preserved visual shape processing (associative agnosia; Farah 2004, chap. 6). Thus, while an apperceptive agnosic might not be able to recognize a gorilla as a gorilla because she cannot resolve its basic features, an associative agnosic might not recognize a gorilla as a gorilla, even though she might be able to accurately draw it.

Given that the inability to resolve basic visual features in apperceptive agnosia might suffice for a kind of blindness to the object, inattentional apperceptive agnosia might suffice for inattentional object blindness. Thus, let us look at inattentional associative agnosia where visual processing can lead to a full-blown object representation, but where the subject is unable to categorize the object. That inattention might lead to such a deficit is not an outlandish idea. Recall the discussion of early and late selection and of load theory in Chapter 1. The idea is that the attentional filter can impose itself at different points in perceptual processing in a way that is sensitive to perceptual load, i.e., the demands made on perceptual processing. One plausible idea will be that in attentionally demanding tasks, conditions are high-load, and one view holds that this pushes attentional selection to earlier parts of processing. To speculate, unattended objects will not be processed to the level of perceptual categorization. Inattentional associative agnosia might be the result.

In the gorilla experiment, while one sees the gorilla, one fails to see it as a gorilla. The task of following a basketball in a cluttered scene of moving white and black objects might be a high-perceptual-load task drawing processing resources away from unattended stimuli (see Cartwright-Finch and Lavie 2007 for correlation between load and inattentional blindness). There might be sufficient resources to process unattended stimuli, but only to a certain level, say, form and color, but not category. If so, then while one might see what is in fact a gorilla, one would not be able to see it as a gorilla. One would not then report seeing something odd, a gorilla among people. Rather, one would simply see the gorilla as one black shape among several other black shapes. One is then conscious of the gorilla as a black object, but not as a gorilla. If so, inattentional associative agnosia is consistent with visual awareness of the gorilla, but also explains the failure to report seeing a gorilla.

One question is whether this approach can explain the Mack and Rock paradigms where more simple stimuli were often used. In any event, let us note that it is a second potential explanation of the subject's failure to report. It is important to see, however, that inattentional agnosia is consistent with the common-sense model and allows for consciousness outside of attention. The idea is that inattention makes you blind to a gorilla's being a gorilla but allows for awareness of other features of the gorilla.

5.6.3 Inattentional apraxia

The logic of inattentional blindness experiments, such as in the case of the gorilla, requires that one abolish attention to a critical stimulus so that the putative gate to consciousness is completely closed. That is, in order to

support the claim that attention gates for consciousness, the experimenter must show that when the gate is closed, consciousness is thereby limited. A successful inattentional blindness experiment requires that the attentional gate is completely closed. At this point, it is crucial to specify what attention is supposed to be. Consider attention as selection for action. The selection for action theorist will note that, to report X, one needs to select X for report. Accordingly, any paradigm that is successful in completely siphoning attention away from X will abolish the subject's ability to report X. But then satisfying a necessary condition for a successful gatekeeping experiment, viz., that the subject not attend to a critical stimulus X, is sufficient to generate the central result: the absence of report of the unattended stimulus X. Notice, however, that this failure to report is independent of whether or not the subject is phenomenally conscious of X. If the subject were conscious of X, the subject will still fail to report X if attention to X, and hence selection for action of X, is prevented. I have derived this consequence by construing attention as selection for action, but one can derive the same result by using a version of the empirical sufficient condition (Chapter 1): if S selects X for report, then S attends to X. It follows that where S is not attending to X due to having attention fully deployed elsewhere, S is not selecting X for report. The same consequences then follow, namely that abrogating attention guarantees failure of report (see also Stazicker 2011, Section 2).

Let us consider the two moments when one might expect subjects to report on a salient stimulus, namely: (a) during the experiment, when the stimulus appears ("Hey, what's that gorilla doing there?"); or (b) retrospectively, when queried after stimulus presentation ("Oh, there was something funny ... I think it was a gorilla!"). In the first case, if attention is fully siphoned away from the gorilla, then one expects that the subject will not report the presence of the stimulus. There will be no selection of the gorilla for report. In the second case, one should think about what selection for retrospective report involves. If I ask you to remember the first five cards I reveal in rapidly flipping through a standard deck of cards and then to later report them, you will select those five cards for memory, and then report on them by drawing on memory. Retrospective report is a memory-guided task, and so selection for such reports requires selection for memory as a crucial component of selecting for action. So, in selecting an object for later report, selection for memory-guided action requires that when the stimulus is presented, it is selected for memory. If not, then the stimulus will be lost to later action (consider change blindness here). But if at the time of stimulus presentation, the capacity for attention is fully deployed elsewhere, then the stimulus cannot be selected for memory so as to serve later action. That is, inattentional amnesia obtains. This suggests that, again, where attention is completely diverted away from the critical stimulus, the possibility of report is abolished.

There is a serious, but unnoted, challenge here to those who want to draw the conclusion of inattentional blindness from the experiments that depend on failure to report. The primary evidence for blindness is failure of report, and yet, when fulfilled, the experimental conditions in all experiments supporting inattentional blindness guarantee that the subject will not report the presence of the stimulus, even if the subject is conscious of the stimulus. Rather than demonstrating inattentional blindness, the experiments, when successful, demonstrate inattentional *apraxia*, an inability to respond to the relevant stimulus precisely because attention is tied up with another task. If this is right, we have as yet no evidence that attention is a gatekeeper for consciousness. Worse yet, we might not be able to gather such evidence, at least with the experimental paradigms considered thus far.⁷

5.7 Consciousness without attention?

Let us turn the tables now and question opponents of gatekeeping, for they claim that there is consciousness outside of attention. Can they demonstrate this?

5.7.1 On gist

Arien Mack has noted that "the claim that gist perception is possible without attention is not trivial. It is the main evidence given for attention-free awareness" (Mack and Clarke 2012, 303). Perception of gist is taken to be the best counterexample to the gatekeeping view, but what is gist? Aude Oliva (2005) has characterized it as a

spatial representation of the outside world that is rich enough to grasp the meaning of the scene, recognizing a few objects and other salient information in the image, which includes all levels of processing, from low level features (e.g. colour, spatial frequencies) to intermediate image properties (e.g. surface, volume) and high level information (e.g. objects, activation of semantic knowledge).

In her terminology, conceptual gist involves semantic information, while perceptual gist involves visual representation. She points out that, in less than 100 milliseconds (ms), observers can recognize the basic category of a scene, as well as aspects of its spatial layout and its global structure. Thus, if you present a picture to me, where it is flashed for only 100ms, I might be able to grasp that it is a street scene or that there are vehicles in the picture, even if I can't report more specific details. Colloquially, I get the gist.

In the current context, the claim is that one can be conscious of gist in the absence of attention to it. But why think this? The work of Mack and Rock (1998, 165–68) is often cited as providing initial evidence that gist is immune to inattentional blindness. In a critical trial of their inattentional blindness paradigm, Mack and Rock superimposed the cross used in their discrimination task over a photo of a breakfast scene or of a scene of two individuals petting a dog with a 200 ms presentation of the image. Performance in the cross-discrimination task was as in their previous studies, but Mack and Rock noted that no subject experienced inattentional blindness in respect of the "essence or gist of the picture" (167). Most subjects could report the gist, as well as certain details.

Is gist immune to inattentional blindness? Some claim that the Mack and Rock result provides evidence that gist is immune (Cohen, Alvarez, and Nakayama 2011, 1166). In commenting on their result, however, Mack and Rock noted that "it is the size of the scene that attracts attention and that once this occurs, attention may be distributed over the entire array, or perhaps those parts of the array that carry most meaning are given priority in processing and thus are more likely to be consciously perceived" (169; my emphasis). Thus, they do not claim that a subject perceives gist without any attention, for bottom-up attention can be captured by gist. Accordingly, Mack and Rock's result is consistent with the claim that attention serves as a gatekeeper, for, in this case, bottom-up attention guarantees that the gist is consciously experienced. Indeed, in their concluding chapter of the book, they remark: "If attention is necessary for perception, then an object presented under conditions of inattention necessarily must capture attention to be perceived" (228), here presumably emphasizing bottom-up attention. Mack has recently noted with some bemusement that the received view seems to be that her earlier work with Rock argues that gist is immune to inattentional blindness, despite their clear claim otherwise (Mack and Clark 2012). That there is such an interpretive disconnect highlights the challenge of keeping track of different gatekeeping conditionals, in particular, here, the difference between top-down versus bottom-up attention conditionals.

Indeed, this remains a constant challenge, for much recent work has focused on whether gist is genuinely immune to top-down inattentional

blindness (Li et al. 2002; Mack and Clarke 2012; Cohen, Alvarez, and Nakayama 2011). But given Mack and Rock's emphasis on gist-triggering bottom-up attention, they can allow that top-down inattention does not lead to gist blindness. Establishing that gist is immune to top-down inattentional blindness does not provide evidence against the general gatekeeping view, even if it provides evidence against a restricted version of it, namely, that consciousness of gist implies top-down attention to gist. Recent work then seems to be at cross purposes with the original interpretation Mack and Rock provided of gist. Indeed, should immunity from top-down inattentional blindness surprise us? If the fire alarm were to go off suddenly while subjects were fully engaging top-down attention in a psychological experiment (which would be true if the experiment was successful), would one expect them to notice it? Isn't the point of bottom-up attention as a circuit breaker that it can make us aware of a stimulus even if attention is fully engaged with current goals (see Chapter 1, Section 7)? Certainly, showing this is an empirical matter, not to be decided by such intuitions, but having a bottom-up-driven circuit breaker does seem plausible.

Despite the issue of being at cross purposes, let us consider attempts to empirically demonstrate or refute top-down inattentional blindness to gist. The general logic of such experiments is to identify a task that plausibly consumes all top-down attention, and then show that the subject is still aware of gist as revealed by report. A constant challenge is how to demonstrate that top-down attention is fully engaged and not available for the detection of gist (for questions about whether this is possible, see Cohen et al. 2012). Work from Christof Koch's group (Li et al. 2002) has been the topic of some recent discussion. In their experiments, top-down attention is pulled away by a demanding task, namely, the subject must report whether briefly presented letters in a group at fixation are the same, or whether some letters are different. At the same time, subjects were presented with briefly flashed stimuli of scenes in the periphery that did or did not contain animals. Subjects were further tasked with making judgments about whether an animal was present or not (this was the gist task). In fact, both tasks were highly practiced, with over 12,000 practice runs. In a "dual task" experiment where subjects had to perform both the letter and animal judgment, Li et al. demonstrate that both tasks can be performed without decrement relative to their being performed individually. They conclude that there is awareness of the gist without top-down attention to it, the assumption being that top-down attention is fully engaged in the letter task.

Against this claim, Cohen et al. (2011) and Mack and Clarke (2012) argue that, in their experiments, subjects do exhibit inattentional blindness for gist during attentionally demanding tasks (Cohen et al. use the same detection task as used in Li et al.). To this, one can raise an earlier worry, namely, that if these experiments succeed in tying up attention away from the gist, then one should expect subjects not to report on gist. The result does not support inattentional blindness over inattentional apraxia. Furthermore, Cohen et al. and Mack and Clark question whether the Li et al. task fully engages top-down attention (but see Li et al. op. cit. for controls that argue that their task does, pp. 95–99). Koch often reports his results as demonstrating awareness in the "near absence" of top-down attention. Of course, there can be versions of the gatekeeping view that claim that consciousness requires more than the near absence of attention (you are conscious of X only if you are attending more than a little to X). Still, in respect of the general gatekeeping view, which is probably what most gatekeeping theorists intend, if there remains even some residual attention, then one will not have a counterexample to gatekeeping.⁸ In fact, one might argue that it is not just that there is some residual attention in the Li et al. experiments, but that there is quite a bit of it. After all, one of the tasks is precisely to report on gist, namely, the peripheral display. The question then is this: How can that task by its very nature fail to engage top-down attention? It is a task that subjects intend to do. Of course, given practice, many features of task performance will be automatized, but, fundamentally, subjects take on the goal of making reports on gist as part of agreeing to do the task (see the definitions in Chapter 1, Section 7). Given the empirical sufficient condition, selection of the gist for report is attention to the gist, and indeed top-down attention to it. This work illustrates why careful formulation of gatekeeper claims is necessary for appropriate engagement (see Section 5.2). What we have seen is some talk at cross-purposes, but also claims that do not fully engage with the general gatekeeping hypothesis. Accordingly, the conceptual distinctions made in discussing the gatekeeping conditionals really must come to the fore in further discussions of gatekeeping.

5.7.2 On crowding and the grain of attention

We are left with the challenge of showing consciousness outside of attention. Let us close with an ingenious proposal by Ned Block (2012) that provides a different way of attacking the gatekeeping thesis. I shall argue that, as it
stands, Block's proposal does not clearly succeed, but that there are many open empirical questions such that, resolved in a certain way, he might be correct.

Block focuses on the phenomenon of crowding. Consider Figure 5.2:



Figure 5.2 A demonstration of crowding redrawn from Intriligator and Cavanagh (2001). They note (figure 1, p. 172): "A simple demonstration which shows the difference between visual acuity and attention. While fixating the cross, the lines to the right are easily seen—they are thin, vertical, parallel, evenly spaced, black, and all about the same height. However, while still fixating the cross, it is difficult or impossible to attend to an individual line in the middle of the group, say, the fourth line from fixation."

Focus your eye on the cross and notice that it is hard to see and attend to each line with sufficient resolution to separate them. The flanking objects crowd the middle object. This makes it seemingly impossible to count each line individually. Furthermore, counting seems to depend on attention to each line, but one is unable to count each line even when trying to attend to the line. This suggests that there are limits to attention's resolving power. Attention is insufficiently fine-grained to allow it to individuate the lines. You can see this by noting the difficulty of following instructions such as: "start from the left, now go five lines to the right, two to the left, four to the right ... " (Intriligator and Cavanagh 2001). Using experimental tasks such as these, Patrick Cavanagh and coworkers argued that the resolution of attention is not as fine as the resolution (acuity) of vision. If so, one might think that one can see each individual line even if one can't attend to a specific one. This would present a counterexample to the gatekeeper view: it is possible to visually experience a line without attending to it.⁹

For argument's sake, I allow that the subject cannot attend to a crowded line in the center of the lines. Interestingly, the basic inference seems to be that individuating an object in a counting task or in some other task is a necessary condition on attending to the object. There might, however, be a simpler reason why attention cannot lock onto an object, namely, that the visual system fails to generate an adequate representation of the objects so

╋



Figure 5.3 Fixate on the + and cast attention covertly to the middle figure. Can you see the middle A? Block refers to this as identity crowding, noting that we can just make out the A.

that one can move from one to the other. This is the point I shall now pursue: Is Block right that one can see the object despite not attending to it?

Block focuses on a case that he calls identity crowding, as exemplified in Figure 5.3 where the central **A** is crowded by the flanking **A**s (it is worth emphasizing that this is Block's terminology, not one that is deployed by empirical theorists of crowding).

Block suggests that we can discern the crowded **A**, identify it with some confidence as an **A**. He elaborates:

Since identity-crowding allows detection (i.e., distinguishing between presence and absence), differentiation from the background, discrimination from other items and visual identification of the items—all consciously—it is difficult to see a rationale for denying that one can consciously see them.

(175)

But in order to see an object, one's visual system must construct representations of an object. If crowding disrupts the formation of object representations, then there will be no visual experiencing of an object. Indeed, the bulk of current psychological explanations of crowding hypothesize that crowding is the result of deficiencies in *feature integration*, a necessary step that is prior to the formation of coherent object representations when features get bound. These explanations are often informed by a simple two-step model of object awareness that involves: (a) feature integration; and then (b) object recognition (Pelli and Tillman 2008; Whitney and Levi 2011). Pelli and Tillman (2008) suggest that "crowding' occurs when objects are too close together and features from several objects are combined into a jumbled percept" (1129), and they conclude that "there is a growing consensus that crowding is the combining of features over an inappropriately large area" (1134). Finally, Petrov and Popple (2007) canvas a variety of possible mechanisms for crowding, the bulk of which concern integration or pooling of features. This suggests a different explanation of Block's identity crowding, namely, that attention selects not an object, but a region and the features within that region. Where there is recognition of a sort of uniformity in features (as in a texture), then, given that one can often see the flankers, one can infer that the central object is the same kind, namely, an **A** (Taylor 2013).¹⁰

Block maintains that there is, in the special case of identity crowding, visual experience of an object despite the failure of attention to it. His main evidence for this is an interesting experiment performed by Jeremy Freeman and Denis Pelli (2007). Freeman and Pelli used a spatial cue in crowded and uncrowded letter displays to aid change detection. Subjects were presented with a letter array, either a circle of well-spaced letters surrounding a fixation point where each letter can be clearly identified (uncrowded condition), or a sequence of closely packed letters in the upper visual field above fixation (crowded condition). For each condition, after presentation of the first letter array, a "post cue" appears on a blank screen at the spatial location of one of the letters. This is followed by the presentation of a second letter array that is either identical to the first array or has a change at the cued location. The task is change detection: Did one of the letters change? What Freeman and Pelli discovered was that, with the cue, subjects could detect a change in crowded and uncrowded displays to the same performance level. In other words, crowding did not have a detrimental effect on change detection when there was a spatial cue (it did have a detrimental effect on change detection without the cue). Block gives two reasons why this result supports his claim of visual object awareness of crowded letters:

First, features must be bound to a location in order to be locationally cued as in this experiment. Second, since many of the pairs of letters overlap a lot in features, the fact that performance does not depend on whether the letters are crowded or not suggests that the features are structured.

(180)

Let us raise some questions for Block. First, Block identified identity crowding as the best case of object perception during crowding, but the Freeman and Pelli experiment did not use identity crowding. Given the randomized choice of eight target letters from 19 possible letters (Freeman and Pelli 2007, 3), most of the crowded displays used in each trial likely did not exemplify identification crowding. Does Block claim object perception in non-identity crowding cases? This would be against the grain of most models of crowding. Given that, in other contexts, such crowding has led theorists to postulate disrupted object awareness, one might expect the same disruption in the Freeman and Pelli experiments. If so, then one needs to take seriously that it is awareness of features, not objects, that explains task performance.¹¹

Second, Block's first point about features bound to a location seems compatible with a feature-based explanation of change detection. That is, so long as a feature is bound to a location, a location cue allows feature attention to select the relevant feature at that location and to compare it with a relevant feature at the same location in the second display, in the absence of the visual system generating a visual representation of an object. So, while object awareness aids performance in the uncrowded condition, it is feature awareness that aids performance in the crowded condition.

Finally, it might be the case that crowding does in fact disrupt basic neural mechanisms needed for the visual representation of objects. There is currently no published work that shows this, but I raise it as an open question. If the normal activity of neurons that serve object representations is disrupted under conditions of crowding, and the normal activity of those neurons is necessary for the awareness of objects (cf. visual agnosia discussed earlier), then crowding might induce a form of visual agnosia. So, for all that one can say from the psychological data, the neural data might in the future speak against Block's interpretation of identity crowding. I suspect that Block will agree that there are open empirical questions. Block's appeal to crowding is to find a different way to leverage the empirical issues to assess the status of the gatekeeping view. These are the type of approaches that we will need if we are to decide whether the gatekeeping view is correct.

5.8 Summary

There have been a number of striking experiments that suggest an intimate relation between consciousness and attention, and this has led some to claim that attention is a gatekeeper for consciousness. We have been debating these issues in philosophy and science for some time now, and there seems to be a sense that the issues have been settled to a large extent, or that the topic is stale. Far from it! If the discussion of this chapter is correct, all the central questions are yet to be answered, including:

- What is meant by "blindness" in inattentional blindness?
- What notion of attention is in play?
- Which of the many gatekeeping theses are proponents of inattentional blindness defending?
- How can we garner experimental evidence in favor of or against the gatekeeping hypothesis?

In vision, we capture the gatekeeping model by emphasizing its implication of inattentional blindness: when attention is drawn away from X, we do not visually experience X. The strongest version of the gatekeeping view entails that attention defines the boundaries of the visual field, dividing that which we can see from that to which we are blind. This form of inattentional blindness emphasizes literal blindness, but the evidence that we have discussed in this chapter is insufficient to establish it. There are weaker versions of inattentional blindness such as inattentional agnosia where attention does not define the boundaries of the visual field but does determine in a different way what is visible. Inattentional agnosia, however, is consistent with the common-sense model, for some aspects of consciousness are not determined by attention. The challenge, then, is to find a different experimental approach that allows us to address gatekeeping and to adequately cash out the basic gatekeeping claims. This issue is far from stale and there is much work to be done.

Suggested reading

Mack and Rock (1998) provide an extensive overview of their work on inattentional blindness. Mole (2008a) provided some early criticisms, while Kentridge et al. (2008) provide a reply. Among cognitive scientists, Koch and Tsuchiya (2007) argue that consciousness and attention are separate processes, though see Mole (2008b) for a reply. Simons and Ambinder (2005) specifically discuss change blindness. Dretske (2007) provides a philosophical counterpoint. Noë (2002) collects essays exploring the empirical and philosophical significance of inattentional blindness.

Notes

¹ For a summary of empirical arguments against gatekeeping, see (Koch and Tsuchiya 2007; Tsuchiya, Block, and Koch 2012). For arguments in favor of gatekeeping, see (Prinz 2012).

2 How about studies with normal subjects? An experiment by Jiang et al. (2006) involves the clever use of interocular suppression: when two distinct images are each simultaneously presented to separate eyes, visual awareness oscillates between the two. That is, rather than seeing both images superimposed on each other, we first see one, then see the other, say, a house, then a face. If one of these stimuli is rendered more "powerful," e.g., by increasing its luminance or by enhancing its contrast, it can essentially suppress awareness of the other stimulus. Let us say that the suppressed stimulus is invisible. Subjects do not report it, so we conclude that they are not visually aware of its presence in the sense that they lack phenomenal experience of it.

Jiang et al. used strongly arousing images of nude men and women as the invisible image. To one eye, the experimenters presented the arousing image and a scrambled version of it, each on opposite sides of a fixation point. To the other eye, a pair of "noise" images was presented in a similar fashion. In the experimental conditions, when confronted with both pairs of images, subjects were visually aware of only the noise images (when they were aware of some difference, those trials were discarded). This presentation was then used for a modified version of the Posner spatial cueing paradigm where the invisible erotic image was used as a direct spatial cue, presented to only one eye. After presentation of the cue for 800ms and a 100ms stimulus onset asynchrony (the gap between cue and target), a Gabor Patch was presented either in the cued location or in an uncued location for 100ms. The Gabor patch was rotated either clockwise or counterclockwise, and subjects were required to report their orientation.

Strikingly, the invisible cue, the nude image, induced a benefit for spatial attention as measured by response accuracy: reports were more accurate in the cued versus uncued condition when the cue was valid. This is a standard spatial attention effect in the Posner paradigm. More striking was the observation that the effects segregated along gender and sexuality. Female nudes captured the attention of heterosexual males, while male nudes (interestingly) repulsed their attention, as based on reaction times. In contrast, male nudes captured the attention. Further results showed that male and female populations could be subdivided along sexual orientation. For example, gay male subjects showed similar tendencies in respect of the images as their heterosexual female counterparts.

Experimentally, the cuing paradigm secures attention to X and the mask obliterates consciousness to X, where X = nude image. So, it looks like we

174 ATTENTION AND INATTENTIONAL BLINDNESS

have a counterexample to conditional (1): that when one is attending to X, one is conscious of X. We might resist the claim, however, for the Posner paradigm is used as a test for *spatial* and not *object* attention. If that is the case, what the Posner paradigm secures is attention to Y, where Y = target location, and the mask obliterates consciousness of X, where X = nude image. In other words, the experiment provides a counterexample to the following conditional:

If you are attending to location L, then you are conscious of object O at L (if there is one).

While the result is interesting, it does not seem to be the one we were after. This is why it is important to be clear about which gatekeeping conditional we are assessing. Still, we can argue that the test does provide a counterexample to (1) (i.e., that attention to X entails consciousness of X). After all, it is a nude *body* that is arousing to subjects and not a spatial location. Similarly, for heterosexual males, images of nude males repelled attention. If this result is genuine, then presumably it has nothing to do with the spatial location that the subjects found automatically repelling, as reflected in reaction-time cost. Rather, it is the object, the nude. One might then argue that in the non-repelling case of an image of a nude, the object is what attracts attention. Indeed, recall that there is a scrambled version of the image in the other half of the visual field, but the benefit was seen where the probe occurred at the location of the nude, and not of its scrambled counterpart. This suggests that we should substitute the same values for X in (1), either object (a nude) or feature (being nude).

- 3 The following text is reproduced from a forthcoming review of mine (Wu, forthcoming) on Prinz (2012). It is reproduced here by permission of Oxford University Press on behalf of the Mind Association.
- 4 For a recent discussion, see also Taylor (2013).
- 5 For an influential model of the attentional blink, see (Chun and Potter 1995). The nature of the task has a strong influence on the attentional blink (AB). When subjects are told of the likely stimulus onset asynchrony (gap) between the two targets, the AB is reduced, suggesting that the nature of the task and the subject's orientation to it has a lot to do with the AB effect. How we deploy attention, namely, to select stimuli for action, is a critical feature in determining AB (a point also made by Awh et al. 2006).
- 6 Those who lose sight later in life might have phenomenal vision in the sense that their initial experiences of sightlessness are like experiences of darkness. But that is a form of visual experience, albeit an impoverished one.

- 7 As I noted earlier, the points raised here also apply to other cases, such as the attentional blink and hemispatial neglect. For in both, if attention is not available to be deployed to the relevant objects, then we do not expect the subject to report them.
- 8 Again, it is worth raising questions about the conception of attention that allows for talk of *residual* attention. Further, note that in the Li et al. experiment, the question of gatekeeping was probed in a dual-task paradigm, namely, that subjects performed a letter-judgment task and an animal-detection task. On the selection for action view, we have two forms of top-down, goal-directed attention, given task instructions. It is hard to understand why this case counts as a form of top-down attention in the central letter-judgment task, but not in the peripheral animal-detection task. The very nature of the task instructions ensures that the selection of stimuli for report is goal-directed and top-down in both cases (given training, we expect some level of automaticity, but, as I argued in Chapter 1, this is consistent with it being top-down). This argues against using the experimental results in support of consciousness in the absence of top-down attention.
- 9 One response given by Michael Tye (2009) denies that the antecedent holds, namely that one is visually conscious of object *O*, say, the middle line. The reason is that Tye also holds that to be visually conscious of *O*, then one must be able to attend to *O*. Given that the debate concerns whether attention or a capacity for attention is a necessary condition on perception, Tye's assumption cannot be drawn on in this context.
- 10 This basic point was raised by Mazviita Chirimuuta during Ned Block's visit to Chris Hill's consciousness seminar at the University of Pittsburgh, 2013, where Block discussed his article. In his article, Block rejects this response and further discusses it in his (Block 2013). For a different discussion of Block's argument, with emphasis on unconscious vision, see (Richards 2013).
- 11 No doubt the Freeman and Pelli experiment will show the same result if identity crowding conditions are used. One response is that the featurebased explanation in the original experiment might apply in the identitycrowding version.

6

ATTENTION AS THE GATEKEEPER FOR CONSCIOUSNESS: COGNITIVE ACCESS

6.1 Introduction

In Chapter 5, I contrasted the commonsense model with the gatekeeping model of the relation between attention and consciousness:



Perceptual Automotive Perceptual Consciousness — Report

Figure 6.1 The Gatekeeper and Common Sense Models.

This chapter continues with a different elaboration of the two models that focuses on a connection between attention and cognitive access. Cognitive access to X -access to X by (some form of) cognition – implies attention. The canonical case of cognitive access to be discussed at length is access to X by working memory systems, where this access is mediated by attention. On some views to be discussed, attention is for working memory, and in

that way attention is for consciousness. This leads to the following elaboration of the previous models:



Figure 6.2 Cognitive Access and the Gatekeeper and Common Sense Models.

The crucial new element is the insertion of working memory as tied to cognitive access, and while there are complications, given how "access" is used, the simplest elaboration is to identify attention as for working memory and access as encoding in working memory. This results in a switch in emphasis in respect of the gatekeeping thesis. In the last chapter, the focus was on what attention is directed at in perception as determining the character of consciousness. In this chapter, the focus is on what attention delivers to working memory as determining the character of consciousness. Put another way, the shift is from focusing on attention in perception to attention for memory.

As the issues regarding memory and consciousness have been discussed in terms of access, Section 6.2 introduces Ned Block's notion of access consciousness, discusses different applications of the notions of access and accessibility, and ties access to a notion of attention for cognition. Section 6.3 examines two empirical theories of consciousness that take attention for working memory as a necessary condition for phenomenal consciousness. Then, Section 6.4 presents a famous experiment by George Sperling that has provided support for those who argue that attention does not limit phenomenology. Section 6.5 introduces Block's thesis that phenomenology overflows access, while section 6.6 discusses different responses to Block's thesis and considers the experimental evidence relevant to assessing that thesis. Finally, Section 6.7 briefly discusses a neurobiological argument for overflow due to Victor Lamme.

6.2 Phenomenology and access

Ned Block (1995) introduced a distinction between phenomenal consciousness (P-consciousness) and access consciousness (A-consciousness). As discussed in previous chapters, phenomenal consciousness is what it is like for the subject. Block characterized A-consciousness as follows:

A state is A-conscious if it is poised for direct control of thought and action. To add more detail, a representation is A-conscious if it is poised for free use in reasoning and for direct "rational" control of action and speech. (The "rational" is meant to rule out the kind of control that obtains in blindsight.)

(In the version printed in Block 2007b, 168)¹

A-conscious representations are poised for access in the sense of being accessible for use by action systems (with "action" broadly construed). When such representations are in fact used, then they are accessed. So, the central notions are access and accessibility. These notions, however, must be deployed with care. One of Block's early examples of P-consciousness without A-consciousness is the following:

Suppose that you are engaged in intense conversation when suddenly at noon you realize that right outside your window, there is – and has been for some time – a pneumatic drill digging up the street. You were aware of the noise all along, one might say, but only at noon are you consciously aware of it. That is, you were P-conscious of the noise all along, but at noon you are both P-conscious and A-conscious of it.

(Block 2007b, p. 174)

How does the access/accessibility distinction apply in this context? Certainly, before one notices the drilling, one doesn't have access to it in the sense that it does not guide or prompt a report of the sound. Were one to make a report, then one accesses that information to guide behavior. Block also points out that you might realize that the drilling has been going on for some time. This realization calls upon information from memory, but until

the memory is recalled, it is only accessible to and not yet accessed by report.² Yet there is a further dimension, for imagine that as the buzzing occurs during your blissful unawareness of it, your perceptual system registers the sound. There might be a further step between registering the sound in perception to being accessible to encoding in memory, so it is possible to have perceptual information about the drill without this being accessible to memory. This leads to two further distinctions: (a) perceptual representations of the drill that are accessible to working memory, but short of being actually accessed by (encoded in) working memory; and (b) perceptual representations that are not even accessible to memory. The challenge is that the notions of access and accessibility can be used to describe different points in processing. This means that talk of cognitive access or accessibility might refer to different things, and that in discussing the elaboration of the two models, one must be explicit about which meaning is intended lest confusion ensue.³

To regiment the use of the notions of access and accessibility, consider the following flow of information:



Figure 6.3 The Flow of Information via Working Memory.

Given discussion in the literature on cognitive access, four stages are salient:

- 1. Perceptually encoded, currently inaccessible, but potentially accessible to memory
- 2. Perceptually encoded, accessible to, but not yet accessed by, memory
- 3. Encoded in memory, accessible to, but not yet accessed by, behavior
- 4. Accessed from memory to guide behavior such as report.

To keep things orderly, I will use the access/accessibility distinction to describe the flow of information from perception to action where one always speaks of X's accessibility-to-Y or X's being accessed-by-Y (i.e., Y's accessing X). In these locutions, Y identifies a system to which information from X is sent (e.g., working memory, reporting systems). So, in (1), perception is only potentially accessible to working memory; in (2), perception is actually accessible to working memory; in (3), perception is in fact accessed by

working memory, but only accessible to behavior systems; in (4), working memory is then accessed by behavior systems. The required regimentation is to always be clear on what the arguments for X and Y are. The focus in what follows is largely on (2) and (3) with ultimate emphasis on (3).⁴

Is A-consciousness necessary for P-consciousness? Note that there are two interpretations of "A", namely, "access" or "accessibility". To disambiguate, I will drop talk of A-consciousness and focus on the difference between access and accessibility. This leads to the following claims:

```
(A1) Subject S is P-conscious of X only if X is accessed by S.
```

```
(A2) Subject S is P-conscious of X only if X is accessible to S.
```

For these claims, that X is accessed or accessible to S implies that X is accessed by, or accessible to, respectively, S's working memory. If access or accessibility is tied to attention, then we have a version of the gatekeeper view. It is prima facie plausible that access is tied to attention, for we access an item X for some T by selecting it for T (e.g. let "T" stand for task, action, or phenomenal consciousness). For the relevant T, say a task, selection of X for T suffices for attention to X for T. Since accessibility is defined as potential access, both notions are then tied to attention. (A1) and (A2) can then be used to derive the following gatekeeping (GK) theses:

 (GK_{A_1}) Subject S is P-conscious of X only if S attends to X for working memory.

(GK_{A2}) Subject S is P-conscious of X only if S could attend to X for working memory.⁵

If one were inclined to think that actual report of a stimulus (verbal or some relevant behavior) is a necessary condition for P-consciousness, then one would endorse (A1) and hold that P-consciousness arises only when processing reaches stage (4). It is not clear that anyone holds this view except, perhaps, a hard-headed behaviorist. Instead, most hold that reports or relevant behavior provide evidence for phenomenal consciousness, but that actual reports are not necessary for consciousness. Those who endorse (A1) will require access in terms of stage (3) as necessary for phenomenal consciousness (see Global Workspace Theory in Section 6.3.1). Encoding in working memory determines the content of consciousness. Those who endorse (A2) will only require accessibility in terms of (2) as necessary for

phenomenal consciousness (see Attended Intermediate Representations Theory in Section 6.3.2). Finally, those who deny that P-consciousness implies access or accessibility will claim that one can have perceptual consciousness without reaching any of stages (2)-(4). To return to (1), there can be perception that is not (currently) accessible to working memory, and yet is conscious. It is not clear that anyone holds this view.⁶ I will focus on (A1) since all the parties in the debate endorse some version of (A2).

6.3 Two empirical theories of consciousness

This section examines two empirical theories of consciousness that provide accounts of access and its relation to phenomenal consciousness: the Global Workspace Theory and the Attended Intermediate Representations (AIR) Theory. The first theory was initially presented by Bernard Baars (1988). although I shall focus on recent elaborations by Stanislas Dehaene and Lionel Naccache (2001); the second is defended by Jesse Prinz (2012). Both theories share an assumption about the relation between information carried by neurons and the contents of consciousness, namely, the content realization principle (CRP):

(CRP) There is a necessary correlation between the content of consciousness and the information carried by the neural realizers of consciousness.

CRP implies that conscious content correlates or covaries with neural information. Thus: let neural population N realize conscious state C with content P. Then the information I in N realizes P such that, where there is a change in the content of C, there is also a change in information in N, and vice versa. CRP leads to the following question: Why does some information rise to the level of conscious content while other information does not?

6.3.1 The Global Workspace Theory

In the prologue to his book, In the Theater of Consciousness, Bernard Baars writes: "Consciousness seems to be the publicity organ of the brain. It is a facility for accessing, disseminating, and exchanging information, and for exercising global coordination and control" (1997, 7). This functional conception that originated with Baars (1988) has been more explicitly linked by Dehaene and Naccache to the organization of the brain:

the human brain also comprises a distributed neural system or "workspace" with long-distance connectivity that can potentially interconnect multiple specialized brain areas in a coordinated, though variable manner ... The global workspace thus provides a common "communication protocol" through which a particularly large potential for the combination of multiple input, output, and internal systems becomes available.

(2001, 13)

We shall focus on Dehaene and Naccache's account.⁷ The general picture can be depicted as follows:



Figure 6.4 Model of the neural global workspace from S. Dehaene, M. Kerszberg and J.-P. Changeaux (1998) "A neuronal model of a global workspace in effortful cognitive tasks." Proceedings of the National Academy, USA 95: 14529–34. Copyright (1998) National Academy of Sciences, U.S.A. Figure courtesy of Stanislas Dehaene.

This "neural" version of Global Workspace Theory focuses on the structure of specific networks in the brain. It is important to note, however, that there is no single brain structure that constitutes the global workspace, though neural workspace theorists tend to emphasize the frontal and parietal lobes (the fronto-parietal network). Rather, the issue concerns the activity of regions of the brain that have the requisite connectivity. Thus there is an

absence of a sharp anatomical delineation of the workspace system. In time, the contours of the workspace fluctuate as different brain circuits are temporarily mobilized, then demobilized. It would therefore be incorrect to identify the workspace, and therefore consciousness, with a fixed set of brain areas. Rather, many brain areas contain workspace neurons with the appropriate long-distance and widespread connectivity, and at any given time only a fraction of these neurons constitute the mobilized workspace. (op. cit., p. 14)

The workspace is then not just an anatomical notion but a functional characterization of a widely distributed, dynamic neural network. It is a network that makes available information to multiple systems. This requires that the network be realized in circuits that have broad and long-range connections to other parts of the brain. Information that is in the workspace can then be broadcast to (accessed by) other systems.

To understand how this provides a theory of consciousness, one must understand what Dehaene and Naccache take consciousness to be. They account for a transitive notion of consciousness, as when one speaks of the consciousness of color (recall CRP). Further, this notion of consciousness is necessarily tied to reportability (indeed, they call transitive consciousness "access to conscious report" (Dehaene et al. 2006).⁸ On their view, the idea of consciousness that is not reportable, i.e. accessible, is empirically empty (Naccache and Dehaene 2007). So, consciousness of X requires that information regarding X be encoded in the workspace so as to be accessible to guide report and other behaviors. This implicates working memory. How then is information from perception that is available to working memory encoded in working memory? Dehaene and Naccache attribute this role to top-down attention. It is attentional selection that determines which accessible perceptual representations become encoded in, and thus accessed by, working memory. In terms of the neural global workspace, the upshot of attention is that a larger part of the workspace becomes active, spanning the parietal and frontal regions. The global workspace is thereby engaged. Recalling CRP, one can say that, on the Dehaene/Naccache theory, subjects are in conscious states with content P when relevant information is modulated by attention such that it is encoded in the global workspace and accessible to behavior. Attention, then, is the gatekeeper for consciousness by serving as the gatekeeper for working memory (recall Figure 6.2).

6.3.2 Attended Intermediate Representations (AIR) Theory

Jesse Prinz (2012) has argued for a theory of consciousness that endorses the following:

(AIR) Consciousness arises when and only when intermediate-level representations are modulated by attention.

While AIR (Attended Intermediate Representations) applies to all forms of consciousness, Prinz largely focuses on visual consciousness. Accordingly, AIR applied to vision holds that attention to intermediate visual representations is necessary and sufficient for consciousness. In earlier chapters, I argued against sufficiency and responded to Prinz's defense, but, in this chapter, it is the necessary condition that matters.

Prinz begins with ideas originally presented by Ray Jackendoff (1987) who himself drew on David Marr's (1982) seminal book, *Vision*, where visual processing is divided into distinct stages. For discussion purposes, understand the division of visual processing as follows:

Low-level vision: where basic features such as edges are processed;

Intermediate-level vision: that represents the world in a viewpoint-dependent way capturing object boundaries, textures, and depth;

High-level vision: that abstracts away from viewpoint and involves categorical representations of objects and properties.

Thus, if the visual stimulus is Bill Clinton's face, then low-level vision encodes basic visual properties like the boundaries of the face; intermediate level vision encodes a viewpoint, say, a lateral profile of the face, if one is looking at Clinton from the side, and high-level vision encodes its being Clinton's face, a representation that might be activated whatever view one has of his face (e.g., head-on versus from the side). Where does visual information become conscious? Prinz follows Jackendoff in emphasizing that consciousness arises at the intermediate-level, for the representational content of visual experience correlates best with intermediate rather than low- or high-level vision (again, recall CRP). That is, visual experience is tied to a viewpoint and in some sense presents objects as relative to the location that the perceiver occupies. Objects look the way they do given that viewpoint (cognitive scientists speak of egocentric representations which are, presumably, at least a subset of the intermediate representations Prinz has in mind). Prinz makes further claims about the neural realization of intermediate representations, suggesting that they involve specific parts of the visual system such as visual areas V2, V3, V4, and V5 (also known as MT, the middle temporal area), among other areas (2012, p. 52).

Of more direct concern is Prinz's conception of attention. His strategy is to look for a common mechanism which is found in all cases of attention and which might then serve as the referent of the term "attention" (2012, p. 91). A good candidate is change in information flow. Specifically, a stimulus that is attended

becomes available for processes that are controlled and deliberative. For example, we can *report* the stimulus that we consciously perceive, we can reason about it, we can keep it in our minds for a while, and we can willfully choose to examine it further.

(92)

Given the discussion in earlier chapters, this passage might make one think of change in information flow as selection for action. Prinz's focus, however, is more specific, for he sees a connection to working memory: "attention can be identified with the processes that allow information to be encoded in working memory" (93). He characterizes working memory as "a shortterm storage capacity that allows for 'executive control'" (92). One might wonder whether this leaves out a simple form of attention, e.g., when one is directly acting on an object currently perceived. Here, attention might seem to serve action, not working memory.⁹ In any case, the link between attention and working memory provides Prinz a functional analysis of the folk concept of attention (95) and leads to an unpacking of AIR:

(AIR) Consciousness arises when and only when intermediate-level representations undergo changes that allow them to become available to working memory.

(97)

Note that both AIR and Global Workspace Theory acknowledge a role for attention and working memory, in that it is attention for working memory that explains conscious content (hence, attention for cognition). AIR theory differs from the Global Workspace theory in that, while the latter ties conscious content to information *encoded* in working memory, AIR ties it to information *available* to working memory. Put in terms of access and accessibility, Global Workspace theory takes P-consciousness to depend on access, in that information must be accessed by working memory (hence, encoded; thesis (A1) Section 6.2); AIR theory takes P-consciousness to depend on accessibility in that information must be accessible to working

memory (thesis (A2), section 6.2). In part, Prinz favors AIR due to some evidence suggesting that the elimination of working memory, and thus working memory encoding, does not eliminate consciousness (2012, Chapter 3).¹⁰ What one can say is that both AIR and Global Workspace theory agree that attention, in the sense of selection that is tied in some way to working memory, is necessary for phenomenal consciousness. Thus, both theories entail gatekeeping.

6.3.3 Attention, A- and P-consciousness: the issues

There are two theses about the dependency of P-consciousness on access/ accessibility.

- (A1) Subject S is P-conscious of X only if X is accessed by S.
- (A2) Subject S is P-conscious of X only if X is accessible to S.

Attention is relevant because it provides a route to cognitive access/accessibility in respect of selection for working memory. Thus, on AIR theory, attention renders intermediate perceptual representations accessible to working memory while in Global Workspace theory (GWT), attention allows working memory to access perceptual representations.

(GWT) X is accessed by S only if S attends to X. (AIR) X is accessible to S only if S attends to X.¹¹

In this way, attention, by being tied to access or accessibility, serves as a gatekeeper for phenomenal consciousness by being a gatekeeper for working memory. For (A1) conjoined with (GWT), and (A2) conjoined with (AIR), imply a familiar gatekeeping conditional: Subject S is P-conscious of X only if S attends to X, i.e., only if S selects X in some way for working memory. Let us simplify matters by focusing on GWT and (A1). (A1) implies (A2) in that if S accesses X, then X is accessible to S.¹² This then makes the challenge to the gatekeeping view very specific: Can it be demonstrated that there is phenomenal consciousness outside of what is encoded in working memory?

In the last chapter, I argued that experiments aimed at teasing apart different models of attention's role in consciousness falter because the conditions of inattention needed to show inattentional blindness suffice to undercut the ability to report the stimulus. At the same time, since report is how one gains access to consciousness and report implicates attention, then it looks like the primary evidence for consciousness cannot also provide evidence for consciousness in the absence of attention. The consciousness one attests to in a report is also consciousness to which one is attentive. This raises what Ned Block (Block 2007b) has called a *Methodological Puzzle*: How can one experimentally address the issues given the limitations just noted?

Block's solution to the puzzle is to deploy inference to the best explanation. That is, he advocates choosing the model that best explains the relevant data. His argument against the gatekeeping view can be reconstructed as follows:

- 1. Visual working memory (the workspace) has a limited capacity.
- 2. Overflow: phenomenology has a higher capacity than working memory
- 3. "The control of working memory is in the front of the head" (496).
- 4. Arguably, the "core neural basis of visual phenomenology is in the back of the head" (ibid.).
- 5. If one assumes that the machinery controlling working memory is necessary for visual phenomenology, then one cannot explain overflow.
- 6. If one assumes that the machinery controlling working memory is not necessary for visual phenomenology, than one can explain overflow.

The idea is that the best explanation of overflow is that the machinery of phenomenology is distinct from the machinery of working memory. Overflow implies that phenomenal consciousness is not limited by attention for working memory, for the capacity of phenomenology is greater than the capacity of working memory. But why accept overflow?

6.4 Sperling, partial reports, and iconic memory

In 1960, George Sperling published a paper titled, "The information available in brief visual presentations" (Sperling 1960). Sperling's question was: How much does one see in a glance? To answer this, he presented visual stimuli to subjects for very brief durations, an experimental paradigm with a very long history dating back to the late nineteenth century. In those earlier studies, subjects were asked to report what they saw of briefly presented stimuli, and thus they had to draw on memory of the stimuli.¹³ As Sperling noted, a repeated early finding was that subjects could only report a subset of what was presented to them. At the same time, subjects typically claimed to see more than they could report. Sperling's advance was to take this sense of seeing more than can be reported as a basis for asking a further question: Does one see more than can be remembered? An answer to this question is directly relevant to the gatekeeper view.

Sperling's goal was to determine the informational capacity of what is seen and whether this is tied to the capacity of memory. He recognized, however, that if memory for report is limited, then attempts to report everything that was seen (total report) can never exceed the capacity of memory for report (what is now called working memory). Accordingly, he opted for a partial report paradigm: the subject reports only on part of what was seen, as determined by task instructions. Sperling's ingenious approach was to use partial reports to circumvent the limits of working memory as revealed in total reports.

7 I V F X L 5 3 B 4 W 7

Figure 6.5 Letter Array in the Sperling Partial Report Paradigm.

Sperling presented subjects with stimuli containing a number of numerals and letters (from 3-12) in various configurations. A sample 12-figure configuration is reproduced here:

When his subjects were asked to give a total report of the identity of the letters flashed, they were able to report on average 4.3 letters (experiment 1, p. 6). This estimate was stable across changes in stimulation durations from 0.015 to 0.5 seconds (experiment 2, p. 6). In his third experiment, Sperling shifted to partial report where subjects were required to report no more than four letters from a stimulus display. Consider then a presentation of 12 letters in three lines arranged top to bottom, with four letters per line (Figure 6.5). Sperling used a tone after stimulus presentation to indicate randomly which line subjects were to report. He assumed that if subjects used the tone to tap into a specific part of a memory representation of the

array, namely, that corresponding to the cued line, then by taking the number of letters reported in partial reports and multiplying it by the number of lines in the array, he could obtain an estimate of the total number of letters that were seen. Doing this, Sperling found the number of letters reported to be on average 9.1, about three of four letters in each line. In other words, using partial reports, what was perceptually available—and presumably seen—was measured to be about nine letters; using total reports, what was remembered was measured to be about four letters. So, visual capacity exceeds working memory capacity. The effect is called the partial report advantage and lasts for about 300ms after the stimulus is removed, the stimulus offset. The work has become one of the classic experiments in modern psychology.

What seems to be largely uncontroversial is that Sperling showed that: (a) what is seen, in the specific sense of information processed by the visual system, can persist after stimulus offset; (b) that it can be accessed in report as in the partial report paradigm; and (c) the content of what is seen exceeds the content of what Sperling spoke of as immediate memory (i.e., working memory). It is a further question how to use Sperling's results to adjudicate questions about the gatekeeper view.

The persistence of what is seen after stimulus offset is visual persistence. Max Coltheart (1980) suggested, however, that "visual persistence" is ambiguous between neural, visible, and informational persistence. By "neural persistence" Coltheart referred to the persistent activity of visual neurons after the stimulus is removed; by "visible persistence" he meant the continued visibility of the stimulus after offset, such as in an afterimage; finally, by "informational persistence", Coltheart intended the continued accessibility of the stimulus after offset, referring to this as iconic memory (the term in this context was introduced by Ulric Neisser 1967). The crucial next question is whether iconic memory, what Sperling uncovered in his partial report paradigm, reflects conscious or unconscious perception. If it reflects conscious perception, then the capacity of phenomenology in iconic memory exceeds the capacity of working memory. This then would provide a counterexample to the gatekeeper view.

6.5 Assessing the phenomenology of overflow

One of the central issues in the debate concerns how to characterize the different forms of visual short-term memory (VSTM) elicited by Sperling's paradigm and others inspired by it. Theorists speak of iconic memory, and

recent work by Victor Lamme and collaborators suggests that there is a second form of VSTM, what they call fragile VSTM (Landman, Spekreijse, and Lamme 2003). The final section will briefly discuss fragile VSTM, but I will focus here on three positions regarding the content of any relevant form of VSTM (and thus Sperling's iconic memory) in respect of arrays like those in Sperling's experiments:

- 1. Unconscious: The information is specific and unconscious (or reflects unconsciousness), but can be brought to consciousness, say by attention.
- 2. Nonspecific: The information is conscious (or reflects consciousness), but in some way it is nonspecific, though it can be rendered more specific due to attention.
- 3. Specific: The information is conscious (or reflects consciousness) and highly specific.

"Specific" indicates that the information regarding each letter identity in memory is sufficient to support report of the identity of each letter when appropriately cued (Sperling's result). For example, **A**s are represented as **A**s. Where information is nonspecific, then identity information is in some way not present or degraded, although this idea needs elaboration. Roughly, **A**s are not represented as **A**s.¹⁴ Talk of "reflects" acknowledges that the memory system itself might not be conscious, though it is a trace of a conscious or unconscious state. For letters in Sperling's array, (1) holds that the relevant representation is unconscious; (2) maintains that the subject consciously perceives the letters, but not necessarily as letters, but rather (perhaps) as symbols, shapes, or even as a jumble of features (recall inattentional agnosia, discussed in the previous chapter); (3) asserts that subjects consciously see the letters as the letters they are, though they cannot report on all of them.

Those who endorse cognitive access, and hence attention for cognition, as necessary for phenomenology often endorse (1). Block's version of overflow endorses a version of (3): what one sees exceeds what one can report, and in a way that allows for rich detail. But what of (2)? On first glance, one might take (2) to be inconsistent with gatekeeping views, but in fact it is consistent with those views. If this is correct, then (3) is the only viable (or at least clear) anti-gatekeeping position. To see why, consider how Sperling's paradigm differs from the inattentional blindness paradigms discussed in Chapter 5. The aim of the latter paradigms is to ensure that subjects deploy their attention in a specific, focused manner

away from a target stimulus (a dancing gorilla, a large scale change in a scene). Yet in Sperling's paradigm, subjects cast perceptual attention in the broadest manner possible, namely to the entirety of the letter array. There is no issue of distraction here. Accordingly, the specific issues regarding gatekeeping that are raised by Sperling's paradigm and similar experiments concern the limits of attention on the output side, irrespective of how much attention might be deployed relative to its targets.

Let me report my own phenomenology when experiencing one of Sperling's displays (12 letters, three lines of four letters).¹⁵ Let's call the cued letters subjects are to report the reported letters and the remaining letters, the unreported letters (this is a bit rough, but it will do). The letters I can report (reported letters) visually appear to me as the letters they are. That's why I can report them! Yet it also seems to me that the letters I cannot report (the unreported letters) are nevertheless visible to me. I see something at those positions although they don't appear to be specific letters. Rather, the unreported letters seem to be a smudge, as if blurrily seen, perhaps not even symbol-like. I am grasping for an adequate description, but I would venture to say that what it is like for me to see the unreported letters is similar to what it is like to see the letters at the edge of this page when I look at the middle of the page (I admit, I worry that this description is theory-ladened). Among undergraduates I have taught who have been presented with Sperling's stimulus, they have spontaneously suggested something more like Sid Kouider's (Kouider et al. 2010) contention that the figures appear as fragments.¹⁶ So, the phenomenology I and some of my students report is consistent with Nonspecific. Block reports that his phenomenology is more in line with Specific.

Conflicts in introspection are often hard to adjudicate, but proponents of Nonspecific and Specific can allow that subjects do see more than the specific letters they report or remember. This was Sperling's starting point, and it identifies a crucial difference between Sperling's paradigm and inattentional blindness paradigms. For unlike the latter, the crucial stimuli in the former are in *some sense* reported. That is, it is not like the case of the gorilla where subjects make no reports at all regarding it. Rather, subjects notice and make reports about all the letters in the array. The difference is in the *specificity* of the report. The point, then, is that subjects do have access to all the letters, but possibly in different degrees (see also Stazicker 2011, Section 3). This is reflected in their reports that rely on working memory. Thus, Nonspecific is prima facie consistent with gatekeeping views. It is not the case that there are any conscious elements that outstrip cognitive access. Subjects report on what they are conscious of, and this is more than the four specific letters they can name. If this is correct, then it is Specific that is needed to refute gatekeeping views. Hence, proponents of overflow must endorse Specific. The core of the overflow hypothesis is that the content of experience outruns the capacity of access in this way: phenomenology is specific in its content in a way that working memory is not.

What is clear is that the invocation of capacity needs to be made more precise. For Sperling, the task was to name the identity of the letters, requiring the coding of specific information regarding identity (an A or a 3). Here, capacity is measured in terms of letter identity, and the consistent result is a limit of about four. Yet subjects also report that there are more letters visible than the four they identify, so this information about the other letters is also cognitively accessible. Subjects thus recall more information than merely four letter identities, and, in another sense, working memory capacity is greater than four. Not greater than four letter identities, of course, but greater in terms of a different notion of information, say, the resolution of uncertainty. Subjects have information not only about letter identity but also about the array. For example, they can accurately report that there are more items than four letters in the array. Perhaps this additional information concerns gist, or perhaps it is more specific. It is, however, additional information about the array over and above letter identity. There is, then, a counting question regarding measuring capacity. This is a fairly technical matter that will have to be set aside, but more work needs to be done here if proponents either of Nonspecific or of Specific are to make clear talk of capacity. Remember, it was the tools of information theory, in allowing for precise quantification of informational capacity, that Broadbent thought to be a big step forward for psychology, a new language (see Chapter 1 and Appendix).

Before focusing on relevant experiments and differing interpretations of them, let's consider two reasons Block has emphasized in favor of Specific. In his (2007b), Block notes:

- 1. Subjects in experiments attest to drawing on specific phenomenology in making their partial reports.
- 2. Denying specific phenomenology suggests that, when subjects have specific phenomenology restricted to the specific letters they report, then there is a shift from unconsciousness or generic phenomenology to specific phenomenology. Subjects should notice a change, but they do not.

The first point Block mentions is one that defenders of overflow often raise (see also Burge 2007), yet it is unclear how much weight one should give to it. The claim is largely anecdotal. For example, Block (2011, 570) notes Bernard Baars' observation that "subjects - and experimenters serving as subjects - continue to insist that they are momentarily conscious of all the elements in the array." Yet Baars seems to be reporting Sperling's own observations here, not independent studies that provide clear empirical support for the first point. Further, someone inclined to endorse Nonspecific can insist accurately that "they are momentarily conscious of all the elements in the array." The difference is whether one is aware of them in a specific or nonspecific way. Thus, what is needed, but currently lacking, is a systematic study of subjects' reports about their phenomenology in partial report paradigms.¹⁷ It is worth emphasizing that subjects in the experiment know that the stimulus array presents letters, or at least are told so. Sperling's original subjects were told what they would see (letters) and underwent many trials with the same kind of letter stimuli. They knew or expected that the other positions they were unable to report contained letters. Thus, even if they were to report seeing each specific letter stimulus as the specific letter stimulus it is, their judgment might be affected by their expectation, rather than being an accurate readout of what perception gives them in each trial. Certainly, this is a potentially confounding factor. De Gardelle et al. (2009) showed that when pseudoletters are substituted in a Sperling letter array, subjects still think they are seeing only letters. They suggest that subjects' confidence in being aware of all the letters is a cognitive illusion (for a response, see Block 2011).

Let us turn to Block's point that if attention were needed for specific phenomenology, then one would notice a shift from unconscious/non-specific to specific phenomenology in respect of the letters attended to. But would subjects notice such change? After all, change blindness studies show that even when a subject focuses attention, they miss substantial changes in a visual scene. Of course, if one focuses attention on the location of the change, then the change is easily seen. Yet this might explain why one would never notice the shift from unconscious/generic to specific phenomenology. In the partial report paradigm, the proposed change is induced by attention's moving to a location, rather than the change as occurring in a location where attention is already present. It is not clear that under such conditions, a change of the sort Block considers would be obvious.¹⁸ Block's main argument, however, is to draw on interpretations of the experiments that provide the best explanation of the data, an inference to the best

explanation. His claim is that Specific provides the overall best explanation of a diverse set of results. Let us then pursue alternative explanations in light of the distinction between Nonspecific and Specific.

6.5.1 Postdiction

Sperling arrived at his estimate of the capacity of iconic memory by summing each partial report across the total number of rows. One might wonder, however, whether summation is appropriate. It would be appropriate if the iconic memory representation is unaffected by any further processing induced by the cue. In particular, the representation must not be affected by attention as induced by the cue. If so, the subject could "read off" the data from a stable iconic memory representation. Ian Phillips (2011a), however, has questioned this *independence* assumption, i.e., the assumption "that a subject's experience of the stimulus in a [partial report] condition is independent of which report is cued because the cue comes only *after* display offset" (386). To show this, Phillips draws on the phenomenon of postdiction.

Consider the sensory processing of two stimuli, A at time t_1 and B at time t_2 where t_1 is prior to t_2 (this formulation allows that A and B can be processed in different sensory modalities). The counterintuitive idea of postdiction is that sensory processing of B can affect one's experience of A. This idea is counterintuitive if one assumes that sensory experience is more atomistic: one first experiences A, and then experiences B, where later experiences, or at least processing of later stimuli, cannot affect earlier experiences. An alternative is that sensory experience is a more complicated function of sensory processing over time. In particular, conscious experience of A might result from the unconscious sensory processing of A and B. Accordingly, sensory experience might result from sensory processing that spans significantly more than an instant. Let us call such effects postdictive.

Phillips provides an overview of various postdictive effects, many involving temporal offsets between the relevant stimuli (e.g., array and cue) similar to those found in Sperling's paradigm. Some of these are multimodal, involving two senses. Consider one striking postdictive effect found in the *sound-induced visual bounce* where two circles ("balls") are depicted on a screen as moving towards each other. When these two circles intersect and then continue to move, there are two experiences subjects report: the circles pass through each other, or the circles bounce off each other. When a sound suggesting collision is played at, or around, the intersection of the circles, the intersection is more likely to be experienced as a bounce. Interestingly, this effect can occur even when the sound comes 200 milli-seconds (ms) after the initial intersection (intuitively, one expects the experience of collision to work only if the sound comes right at the initial intersection, as if the circles were real balls).

Recall that the independence assumption leads us to infer that there is a uniform representation of the letter array that subjects tap into in different ways, depending on which line is cued. If this representation is or reflects phenomenal visual states, then the capacity of visual consciousness exceeds working memory capacity. Since any of the letters that the subject can report are represented in specific detail, the iconic memory representation represents each in specific detail, as per Specific. Thus, there is a rich phenomenal representation that exceeds cognitive access.

An alternative interpretation invokes postdiction. The content and structure of the underlying representation depends on which line is cued and, hence, on attention. Attention, on this picture, alters the underlying representation that serves task demands. Where the cue directs subjects to the top line, the underlying representation is brought by attention to be in a format that best serves reporting the top line; where the cue directs subjects to the middle line, the representation is brought by attention to be in a format that best serves reporting the middle line, etc. Attention can either bring the targeted line into consciousness from an unconscious representation or it can sharpen nonspecific representations into specific ones. Either way, there is no uniform phenomenal representation underlying reports across conditions. Rather, the nature of the iconic memory representation varies with attention in light of a subsequent cue. Given postdiction, Nonspecific or Unconscious might be the correct account of iconic memory.

6.5.2 Generic representations and the determinable/determinacy distinction

The cogency of Nonspecific depends on how one understands nonspecific representations of letters. Rick Grush (2007), in his commentary on Block (2007a), suggested that the relevant visual representations are as of generic letters. His example concerns how experience represents text on a page at the periphery of the visual field. See now for yourself. Focus on a word at the center of this page, and covertly attend to words at the periphery. Grush is certainly correct that the way they appear differs from the way the words

at fixation appear in the glory of their specificity.¹⁹ Yet how to understand a visual representation of a generic letter is not completely clear. Given my example about objects in the periphery, we might understand the proposal in terms of visual spatial resolution. James Stazicker (2011) has recently developed a line of response to Block, emphasizing limits on visual spatial resolution. Stazicker puts this in terms of the determinate/determinable relationship.²⁰ A standard example of this relation involves colors, say the determinable red and its determinates, crimson and burgundy. Determinates are ways of instantiating the determinable. As Stazicker comments: "To represent something indeterminately ... is to represent it as instantiating a determinable property, without commitment as to which determination of that determinable it instantiates. Roughly, property A determines property B where to have A is to have B in a specific way" (170). So, if a visual representation represents an object as (say) crimson, then it represents the object as being red in a specific way, namely, as crimson.

So, one can appeal to the spatial resolution of the visual system as a constraint on the determinacy of its spatial representations. Since visually representing shapes is a form of spatial representation, the spatial resolution of vision will provide constraints on the visual representation of shape. This idea can be spelled out by understanding, at least in a sketchy way, spatial processing in vision. Think of the retina as containing spatial filters that are sensitive to specific spatial frequencies. Consider a band of alternative black and white lines as in Figure 6.6:



Figure 6.6 Figure showing increasing spatial frequency from left to right. Note that contrast increases from top to bottom, and higher contrasts are needed to adequately see higher spatial frequencies. This is why the lines appear to be taller as one proceeds to the right. Reprinted from G. M. Boynton (2005) "Contrast Gain in the Brain." Neuron (47): 476–77 with permission from Elsevier.

As you can see, the frequency of alternation of the lines increases per unit distance as one goes from left to right. In principle, this spatial frequency can be represented as a sinusoidal wave (represented as cycles per visual degree (*cpd*); your thumb held at arms length covers about two visual degrees from left to right). Where the relevant visual spatial filters can detect high spatial frequency, they can more finely resolve spatial properties such as the gap between two lines. For high-resolution spatial filters, two lines with a small gap separating them can be distinguished; for lowresolution spatial filters, the two lines cannot be distinguished, and the visual system will fail to detect the gap. Spatial resolution is greatest at the fovea and falls off rapidly. Stazicker's emphasis on spatial resolution is an important addition to the debate, but what should one say about nonspecific representations as postulated by Nonspecific?

A natural thought is that the phenomenal upshot of degrees of spatial resolution is degrees of sharpness in visual representation. One way visual experience can be less sharp is for experience to involve blurriness. Let us understand this in respect of the visual experience of the boundaries of a line with a sharp edge (so there is, objectively, no blurring at the edge). An ideal visual system not limited by spatial resolution can represent the edge of the line as at a determinate location, say at y (how one specifies y does not matter beyond it involving a magnitude reflecting position in some appropriate spatial coordinate system). A less determinate representation of the location of the edge might place it within a range, say between x and z, where x < y < z. One can understand this difference in terms of the uncertainty tied to visual information in respect of where the edge is located. Recalling Shannon information theory (see Appendix), one can say that visual information leaves more uncertainty in the second case regarding the location of the line, but resolves it in the first case. This characterization of differences in spatial resolution allows us to speak of determinates and determinables if one wishes: being at y is a way of being between x and z. Moreover, both representations of the location of the edge can be veridical.

But how does this help with specifying a less determinate representation of a letter? Notice that the previous point was a way of spelling out blurriness in terms of the representation of edge location. Now, a letter such as E consists of lines (edges), and to the extent that the letter is experienced blurrily, then the visual information one gains about the structure of the letter, including its edges, carries a high level of uncertainty, first in terms of the location of edges, but correspondingly in terms of the figure

constituted by the edges. It might then be useful to think of spatial resolution in terms of uncertainty. This is a description of spatial content, though it leaves open how best to characterize the corresponding phenomenology. Still, it gives us a handle on how experience should be characterized in *Nonspecific*: the experience of letters is nonspecific in that it is tied to a high degree of uncertainty, indeed not only about spatial information, but also information about other features.

Emphasizing spatial resolution as characterizing Nonspecific, we can raise a question for its proponents. Why can't proponents of Specific, such as Block, acknowledge that the visual system faces limits of spatial resolution, but plausibly insist that these limits are not at issue in the Sperling experiments? All parties agree that in order to explain Sperling's results, the specific identity of the letters must be visually represented *somewhere* in the cognitive system. Accordingly, the spatial resolution of the retinal locations stimulated by the unreported letters must be sufficiently sensitive to allow for determinate short term memory representations of letter identity, for any of these letters can be accessed for report when cued in Sperling's partial report paradigm.

Now the question is this: If one agrees that spatial resolution of the relevant letters is enough to perform the task, then why does the machinery of phenomenology seem to blur those letters in order to generate nonspecific representations as required by Nonspecific? This might seem like a pointless step, like purposely defocusing modern auto-focusing cameras while taking a picture. In terms of information, the idea is that the machinery of phenomenology adds noise to the system, increasing uncertainty. But why not just maintain, at the level of consciousness, the spatial resolution that is already present in iconic memory? If this is correct, why not then take the blurring to be an unnecessary further step such that explanatory parsimony pushes us to accept Block's alternative instead, namely, that the phenomenology is as Specific claims? One might respond by saying that moving from unconscious to conscious representations will inevitably involve a loss of information since it is an extra step in transmission of information, and this loss of information can precisely result in a phenomenology more like Nonspecific.²¹

Perhaps one way to settle this debate is to understand how much information is lost as it moves from step to step in visual processing (say from iconic memory to what comes next in the processing hierarchy). For example, given Sperling's result with cueing, there is good evidence that the identity of many letters is registered by the visual system where this

iconic memory exceeds working memory capacity (nine versus four letters in Sperling's estimation). It also seems that subjects don't just visually experience four letters. Rather, they see more letters but, minimally, only four of those letters as the letters they are. The issue then concerns their experience of the additional letters. In principle, one question that might be raised concerns the decay of information as it is processed and transmitted during visual processing. The idea of decay is that there is an increase in uncertainty about the layout of the array. Information is thereby lost. If that decay is rapid, then one could make the argument that by the time processing occurs that is necessary for visual consciousness, there is insufficient information content to support the detailed phenomenology that proponents of Specific claim there to be. Attention counteracts this loss by helping to maintain some subset of the information content about the letters from decaying when the subject is appropriately cued. Attention thereby preserves information by selecting it for memory, and this forms the basis of the subject's reports. On this view, those letters not selected for working memory cannot be seen in detail because information regarding them is quickly lost. If so, one can question whether consciousness can reflect the detail proponents of Specific aver, and instead argue that the information content present could only support Nonspecific phenomenology. On the other hand, there might be a rate of decay of information regarding the letters that (a) both explains the specific performance Sperling observed, but (b) also allows that the information regarding the identity of more than four letters is preserved at later stages of processing, even if this information is not funneled into working memory. If informational detail remains at later stages of processing, one might have the basis of an argument for Specific. This proposal is admittedly sketchy, but the point is that more detailed models are needed to connect with the behavioral data that has largely driven this debate. We need an alternative approach to pry the two models at issue apart, one that returns to the concrete specification of capacity limits that information theory can provide.

There seems to be a general sense among theorists in this area that proponents of Specific face a steep uphill battle, but let me raise a question for those who endorse gatekeeping and, specifically, the idea that consciousness is limited by cognitive access and attention: What does it mean to say that consciousness is limited by working memory capacity? When discussing Sperling, I spoke of working memory capacity as about four letters plus perhaps gist, but that is not a theoretically useful way to measure information. To explore the issue further, imagine looking at the ocean

from a boat, marveling at the blue expanse that extends to the horizon. Gatekeeping claims that what you experience, a seemingly vast colored expanse, is in some way limited by cognitive access. But how is phenomenal consciousness of a large spatial area limited by working memory? Is working memory essential to one's online experience of the ocean blue? Might the phenomenology of experience of a colored space outstrip working memory? If not, why not? It is hard to understand the awareness of the blue expanse as constrained by working memory. The point is that gatekeeper theorists can't sit at the sidelines, enjoying the spectacle of their opponents climbing a steep hill. Gatekeeper theorists also have a difficult job to do, namely, to provide a concrete explanation of precisely what it means to say that conscious experience is limited by the capacity of cognitive access. As I noted earlier, this talk of capacity must be made more concrete, and until it is, gatekeeping remains a vague thesis. It does not allow us to say concretely in relevant cases what it means for consciousness to be limited in this way. But being concrete is a way to allow for an adequate assessment of the thesis.

6.6 Fragile visual short-term memory

I now briefly consider a neuroscientific argument for the Overflow thesis and Specific by Victor Lamme. Lamme and his coworkers have empirically isolated a different form of visual short-term memory (VSTM), what they call fragile visual short-term memory. This is a form of short-term memory that is intermediate in capacity between iconic memory probed in Sperling's work and working memory. Lamme has used these results in an argument in support of a version of Specific and to leverage the formulation of new explanatory concepts in this area.

Lamme's work on VSTM is important, extending Sperling's original findings. Adapting a change blindness paradigm, Landman, Spekreijse and Lamme (2003) presented subjects with an array of eight rectangles around a fixation point, with each rectangle oriented either vertically or horizon-tally (see also Sligte, Scholte, and Lamme 2008). This was followed by a presentation of a second array where the orientation of only one of the rectangles was changed. The time interval between the arrays varied from nearly 0.5 seconds to about 1.5 seconds in different experiments. Subjects were also provided a cue either (a) in the first array; (b) during the interstimulus interval; or (c) in the second array. Not surprisingly, when the subject is cued in the first array, they are highly accurate in detecting whether

the cued rectangle changes its orientation in the second array, presumably because the cue allows the subject to attend to that object. Perhaps not surprisingly, subjects are also fairly poor at detecting changes when cued in the second array. The striking result is seen when the cue is presented in the interstimulus interval, because now performance accuracy is surprisingly high, even two seconds after the offset of the first array. Here, the cue seems to enhance performance even after stimulus offset, something Sperling observed as well.

Based on these and other studies, Lamme has argued that there are three forms of VSTM:

- 1. Iconic VSTM
- 2. Fragile VSTM
- 3. Working VSTM.

Lamme sees iconic and fragile VSTM as tied to the phenomena that Sperling characterized, and indeed, Lamme speaks of (1) as retinal iconic memory and of (2) as cortical iconic VSTM to emphasize the areas of the visual system that he takes to subserve each. For example, Lamme takes retinal iconic memory to essentially be the afterimage of the display, something that disappears quickly. Nevertheless, the informational content of iconic memory can survive the end of the afterimage, at which point it becomes cortical or fragile VSTM, fragile because it is easily disrupted by new retinal information. In either case, the capacity of (1) and (2) exceeds that of (3).

These are important extensions of Sperling's work, but how does this help provide a distinctive argument for Specific? Again, the central question is what iconic memory reflects: specific conscious or unconscious information (i.e., Specific or Unconscious). Lamme's argument appears to be as follows:

- 1. There is a high capacity VSTM distinct from working VSTM, namely, iconic (cortical/fragile) VSTM;
- 2. Representations in iconic VSTM exhibit many facets of perceptual organization;
- 3. Conscious representations exhibit perceptual organization;
- 4. Many unconscious representations do not exhibit perceptual organization;
- 5. The most parsimonious explanation is to take iconic VSTM to reflect conscious, and not unconscious, representations.

(reconstructed from Lamme 2010, 210-11)

By perceptual organization, Lamme means features like feature binding, figure-ground segregation, grouping, and organization that allows for

illusions. The first two premises are derived from Lamme's own work on VSTM. Lamme notes that the empirical evidence for (2) is an ongoing project, but that many facets of perceptual organization have been observed for iconic representations. Premise (3) is in a way derived from introspection and cognitive access to experience, while (4) is empirical, derived from what is known about early visual processing, which many deem to be unconscious. The central question, then, is why (5) is correct in taking Specific to be the most parsimonious representation?²² Premises (3) and (4) suggest that certain features of perceptual organization tend to track conscious versus nonconscious processing, but it is not clear that to then associate iconic VSTM with conscious processing amounts to an explanatory parsimonious inference. This move does echo Block's strategy of providing an account that makes best sense of all the data, in response to the methodological puzzle, but it is unclear what parsimony comes to here. Until that is clarified, it is not clear that we have solid grounds to endorse Specific from a neuroscientific perspective.

6.7 Summary

What is it with attention and consciousness? Why is it seemingly so obvious and yet so elusive? I have examined whether attention itself entails phenomenal consciousness, though I argued that it does not (Chapter 4). The past two chapters have considered whether attention has a specific role as gatekeeper for consciousness. This is, as we have seen, a difficult question in that it is hard to find a clear way to empirically engage the issues so as to help us decide between alternative models. Let me summarize some lessons from this and the previous chapter:

- 1. The central contrast is between the common-sense model, where consciousness is not limited by attention, and the gatekeeping model, where consciousness is so limited;
- 2. Theorists must provide clearer formulations of which gatekeeper thesis they are defending;
- 3. Inattentional blindness paradigms, where attention is purposely pulled from a stimulus, cannot provide evidence to settle the issue regarding which model is correct;
- 4. Alternative experimental paradigms or approaches must then be found to test gatekeeping;

- 5. The issue of gatekeeping can be emphasized either from focusing on attention's inputs (as in Chapter 5) or on attention's outputs as discussed here, namely, as being for working memory;
- 6. Sperling's Paradigm and similar approaches provide an alternative approach but the experiments are subject to divergent interpretations;
- 7. The first step to a way out is to develop clearer models about information processing and capacity that can lead to predictions about what experience of unreported targets in a briefly flashed array should be like.

Again, as I noted in the last chapter, these are areas where conceptual clarity and new approaches are needed. In light of the discussion over the past three chapters, there is no doubt that attention plays some important role in consciousness. The question, nevertheless, remains: What is its precise role?

Suggested reading

For an overview of the neural Global Workspace Theory, see Dehaene and Naccache (2001); for an overview of the Attended Intermediate Representation Theory (AIR) see Prinz (2012). Lamme (2010) makes a detailed case for an empirical basis for endorsing the overflow thesis. Block provides an extended presentation of the overflow thesis from an empirical perspective in his (2007b) and a more philosophical perspective in (2008). His (2011) provides a summary of recent work. Phillips (2011b) provides a discussion of Sperling type experiments with emphasis on attention.

Notes

- 1 In his (2007b), Block opts to characterize access in terms of broadcasting in the global workspace (see below on the Global Workspace Theory).
- 2 I am indebted to distinctions drawn by Jesse Prinz (2012). See also Dehaene and Naccache (2001).
- 3 For example, Ned Block focuses on the relation between phenomenology and cognitive *accessibility* in his (2007b) which is geared towards the empirical community, while in his (2008), which is geared towards the philosophical community, he switches to talking about phenomenology as overflowing cognitive *access*. To keep things clear, it is then imperative to be explicit about access/accessibility *to what* (system). I am not saying
that Block is confused about this. Rather, for readers to keep track of the meanings behind invocation of access and accessibility, regimentation is required.

- 4 In his (2007a) discussion of Dehaene and Naccache (2001), Block comments on their division between (I_1) "permanently" inaccessible states, (I_2) states that are *accessible* in that were they to be attended to, they would be accessed by working memory (the global workspace, see Section 6.3.1) and (I_3) states that are accessed by working memory. Block points out two notions of cognitive accessibility, a broad sense that covers (I_2) and (I_3) , and a narrow sense that covers (I_3) . It is cognitive accessibility in the narrow sense that is the focus of Block's discussion. Notice that cognitive access in his terminology is access by systems subsequent to working memory. We are using "cognitive" in a different sense, namely, where it refers in the first instance to working memory.
- 5 (GK_{A2}) has to be rewritten slightly to accommodate Jesse Prinz's AIR view to be discussed in later sections, but the current version will do for now.
- 6 Block is sometimes read as endorsing this extreme view, and I suspect the reason is due to the slipperiness of talk of phenomenology outside of cognitive accessibility. Such talk reasonably suggests to a reader that one means that consciousness is tied to stage (1) and thus not even accessible to working memory. Many find such a view barely coherent. As Block emphasizes in later writings, that is not his claim. The previous conceptual regimentation discussed in the text is crucial for clarity.
- 7 Shanahan and Baars (2007) emphasize that their account of the global workspace does not identify it with working memory, but, rather, as something that gives access to working memory.
- 8 Given the previous regimentation, what they should have said is *accessibility* to conscious report.
- 9 Prinz would presumably respond that even here, attention makes the action-guiding representations accessible to working memory. Fair enough, though that is an empirical question that requires a more concrete specification of what it means to make a representation accessible. It might turn out to be false. Nevertheless, wouldn't the emphasis on working memory lose the forest for the trees in the case imagined, where attention's role seems to be to support action?
- 10 For a critical discussion of Prinz's theory, see (Wu, 2013c). See also (Mole 2013).
- 11 There is a slight complication here that makes terminology fraught with potential peril. It is in fact natural to talk about Prinz as emphasizing

perceptual attention, i.e., attention that influences perceptual representations, while Dehaene and Naccache emphasize cognitive attention, i.e., attention for cognition. The relevant modulations here are all "pointing" towards working memory, even if they occur at different points in processing. Accordingly, I group them together as emphasizing attention that is for cognition.

- 12 This sets aside the central issue that Prinz raises, namely that accessibility is what matters for phenomenal consciousness. This is unfortunate, but the issues will otherwise get overly complicated. Here's why. Ned Block can largely agree with Prinz that phenomenology is always accessible. Where both will disagree is whether accessibility entails attention. Prinz says yes; Block says no (Block 2007b). The debate then centers on whether the property of a representation that renders it accessible is one that is brought about or not by attention. This is an interesting question, but difficult to get a clear handle on. One question we can raise to Prinz is the following: presumably, access also requires attention, but then it looks like attention is involved in two steps, namely, making a visual representation accessible and then, when needed, accessing the visual representation for working memory. But you might wonder if attention ever operates like that. Perhaps attention always just enables access which, of course, implies accessibility. There is no stopping point between accessibility and access once attention gets involved. Still, a fuller discussion is warranted, something that space constraints prevent us from pursuing.
- 13 For a discussion of some of this earlier work and an interesting analysis of the issues, see (Phillips 2011a).
- 14 Again, I find it more helpful to think about information in terms of decreasing uncertainty, so degraded information increases uncertainty. Specific holds that the information content in memory regarding the letter is much less uncertain than what is imputed by Nonspecific. That is, information content resolves uncertainty about letter identity (see Appendix A).
- 15 You can see a version of the Sperling stimulus in an online Ted^X talk by Ian Phillips titled "Swimming against the stream of consciousness" which can be obtained by searching on the internet. The stimulus is presented about 1:30 seconds into the video, which also gives a brief summary of Phillips' account of the experiment, something we discuss in a later section.
- 16 Kouider defends a picture of awareness where the letters that are not identified are given to the subject in fragmentary form. This account is tied to certain assumptions about perceptual processing and our access to it. Specifically, Kouider et al. (2010) hold that perceptual processing involves multiple levels, from basic features to higher order categories

(including gist) such that each level can be independently accessed. Among the perceptual levels are those processing representations of fragments or perhaps parts of the objects present in the visual field. On Kouider's account, we can sometimes grasp the gist of the scene without grasping much detail concerning basic features or objects, or we can focus on a feature and not grasp the gist. This allows for the possibility of what he calls *partial awareness*, awareness that is restricted to some subset of visual processing levels.

On the Kouider view, phenomenal consciousness depends on access, but access can involve all, some, or none of the levels of visual processing. In the case of Sperling's experiments, when subjects claim to see more than they can remember, they are responding to partial awareness, where they have access to low-level representations of the stimuli, namely, letter fragments. Some of these letters, e.g., in cued rows, might be accessed at higher levels, namely, those that present the identity of the letters, while others, say in uncued rows, are accessed as fragments. In this way, the account explains subjects' sense that they see more than can be remembered. Alternatively, when subjects claim to see the specific identities of all the letters, they are under a *cognitive* illusion.

In response, Block (2011) points out that the fragments hypothesis nevertheless suggests that consciousness is rich in content, going beyond the letters that the subjects can explicitly report. He notes that if there is disagreement, it is on "how degraded the specific phenomenology is" (2007, p. 532). So, Kouider and Block can agree that either (2) or (3) is correct, as against (1). Consciousness is not limited to the letters that are reported. Nevertheless, Kouider et al. would emphasize that phenomenology nevertheless does not exceed access. For, on their model, the fragments are in fact accessed, and that is why subjects in Sperling's experiment report that they see more than they could report. Indeed, the sense of seeing more than can be (specifically) reported, demonstrates a kind of access and is thus consistent with gatekeeper views. What remains accessible, even after the capacity to remember the specific identity of particular objects is saturated, is the gist, here the presence of fragmentary forms. Thus, to the extent that Kouider and Block agree, it is that experience is rich in a way inconsistent with (1). They nevertheless continue to disagree about whether (2) or (3) is the correct view. This does suggest the independence between the rich/sparse content distinction and the claim of overflow. Overflow implies that content is rich in the sense that it exceeds cognitive access or accessibility. But deniers of overflow can

also claim that content is in a sense rich, in that it exceeds the specific letters that are reported by subjects in Sperling's experiment.

- 17 Block makes passing reference to an observation of Rogier Landman "that the extent to which subjects evince specific phenomenology may be correlated with how well they do in the experiments [such as those reported in Landman, Spekreijse, and Lamme 2003]" (Block 2007b, 531). This is the sort of evidence that would help buttress overflow, but as far as I know, this observation has never been verified or published by Landman.
- 18 See also (Stazicker 2011, 175–76). For a different response to this issue, see (Phillips 2011b, 215).
- 19 Block characterizes generic phenomenology in terms of existentially quantified content, namely the visual system's representing that there is an array of letters, as opposed to the representation of the identity of each specific letter.
- 20 An influential application of the determinable/determinate distinction to the case of the visual experience of blurriness occurs in Tye (2003).
- 21 Block emphasizes that opponents of Specific endorse the unconscious representation of highly specific visual information, an unconscious icon, but points out that there is no evidence for unconscious iconic memory of the requisite specificity. See his (2011) for a brief discussion.
- 22 Lamme provides a second argument that appeals to *recurrent processing*. Here is a description of the flow of information after a visual signal reaches the brain, one consisting of four stages (see Lamme 2010):

Stage 1: A *superficial* feed forward sweep (FFS) of the signal up the visual hierarchy that does not travel deep into the visual system.

Stage 2: Deep processing of the FFS, where the signal travels the entire sensory hierarchy to motor and prefrontal areas.

Stage 3: *Superficial* recurrent processing involving horizontal and feed*back* connections, of a more local nature.

Stage 4: *Widespread* recurrent processing across the hierarchy (cf. the global workspace).

When the stimulus is removed, iconic memory is associated with Stage 3, while working memory is associated with Stage 4 processing. Thus, encoding in the global workspace is tied to Stage 4. For Lamme's *Recurrent Processing Theory* of consciousness, however, the crucial stage for consciousness begins at Stage 3. Since this is prior to the activation of the global workspace at stage 4, Lamme disagrees that *cognitive access* is necessary for consciousness (he might yet agree with Prinz that *cognitive accessibility* is necessary).

Why does Lamme think that consciousness is tied to Stage 3 as well? In short, it is because recurrent processing looks to be a good neural correlate for consciousness, in part because it is looks to be a good neural correlate for perceptual organization, a critical feature of phenomenal consciousness (Lamme allows that there are open empirical questions here; he is offering a hypothesis). This assumes, as does the argument in the text, that there is some important connection between phenomenal consciousness and perceptual organization. It seems that for Lamme's argument to be compelling, there should be a necessary relation between perceptual organization and phenomenal consciousness. But is there?

7

ATTENTION AND DEMONSTRATIVE THOUGHT

7.1 Introduction

We target the world in different ways: we act on it when we grab an object; we talk about it in referring to things; and we think about it. These ways of targeting the world depend on attending to it. In Chapter 3, I discussed the ways in which action depends on attention. Here, I consider how thinking does as well. The focus will be on a class of thoughts that many theorists take to depend on perceptual attention. These are demonstrative thoughts, those which we typically express in language with demonstrative terms such as "this" and "that". Specifically, I will focus on perception-based demonstrative thoughts about objects. Pick a visible object currently in front of you and ask yourself what its color is, what its weight is, where it came from. Here, you have thoughts about that object, thoughts grounded in your perceptually attending to it. The interest in such thoughts is that they are ubiquitous, seemingly deployed every time one thinks about an object on the basis of currently perceiving it. After all, in looking at an object, the resulting thought is not something complicated like the object to the right of the visual field is red, but simply that (object) is red. The question, then, is this: How are such perception-based demonstrative thoughts about objects made possible by attention?

Section 7.2 highlights different targets of attention and various dependency conditionals that express how one form of attention defined by its target (e.g.,

attention to a feature) might depend on another form of attention (e.g., attention to an object). Section 7.3 then discusses different experimental paradigms used to investigate object attention. The distinction between demonstrative versus descriptive thoughts is introduced in Section 7.4, with the former understood to be conceptually simpler. Several arguments in favor of the necessity of attention for demonstrative thoughts are presented in Section 7.5, while Section 7.6 considers three theories that emphasize conscious attention as necessary. Section 7.7 then raises questions about whether conscious attention does play a necessary role for demonstrative thought.

7.2 Targets of attention

Perceptual attention can be directed at different kinds of targets: space, features, and objects.¹ As noted in Chapter 1, spatial attention can be probed using visual spatial cueing paradigms, while feature attention can be probed in vision with visual search and in audition with dichotic listening focused on properties of competing auditory channels. This chapter examines object attention in vision in more detail to highlight attention that is deployed in demonstrative thought about objects.

For ease of discussion, I will speak of attention to space, feature, and object as different forms of attention, although this should not be taken to imply any stronger claim than that one can individuate episodes of attention by the type of target one attends to. Once all three forms are in view, an important question is how these forms relate to each other. Does one form depend in some way on another? For example, is feature attention a necessary condition for object attention? Thus:

If S attends to an object O, then S attends to some feature F of O.

This dependency conditional might seem initially plausible for visual attention, since one might endorse the following:

If S sees object O, then S sees some feature F of O.

Consider also a parallel in audition: if S hears a sound, then S hears some feature F of the sound (say its pitch, timbre, or intensity). In both cases, it seems that failing to see or hear any features of the object undercuts the ability to see or hear that object. The claim about attention might seem to be entailed by this plausible claim about perception.²

One can also introduce the distinction between perceptual and cognitive attention leading to further dependency conditionals. Consider:

If S cognitively attends to (thinks about) an occurrently perceived object O, then S perceptually attends to O

This is precisely what is claimed for perception-based demonstrative thought about objects: one thinks demonstratively about O only if one perceptually attends to O.

In assessing dependency conditionals, it is important to separate (a) what causes attention from (b) what attention targets. After all, it is possible that X triggers attention to Y (where $Y \neq X$) without X itself being attended to. Consider conjunction visual search where a subject attempts to locate a target defined by two features amid a host of distractors. A common metaphor for attention describes it as a spotlight that serially moves from object to object in an attempt to locate relevant features. Yet, plausibly, that the spotlight moves to a location with a feature or object at it, rather than to an empty location, is due to there being a feature or object at the targeted location. The distinction between causes and targets of attention now raises two possibilities in visual search: (1) that feature/object attention precedes spatial attention; or (2) that the feature/object merely causes spatial attention without itself being attended to. If (1) is true, feature attention is necessary for spatial attention. If (2) is true, then features/ objects are triggers of spatial attention, but need not themselves be targets of attention. Theorists differ on precisely this point in respect of the dependency of object attention on selection of features.

7.3 Objects and attention

This section discusses three experimental paradigms that highlight a role for objects in attention. John Duncan (1984) provided an early approach (see also Treisman, Kahneman, and Burkell 1983). He presented subjects with two distinct, overlapping objects within a region of space where subjects were attending, namely, a line superimposed on a box (Figure 7.1). Each object exemplified two salient features: the line was either dotted or dashed and tilted to the left or to the right; the box was either large or small and had a gap on its left or on its right side. Subjects were tasked with reporting two features present in a given trial, either one from each object, or both from the same object. If attention acts as a spotlight, it should include both objects.



Figure 7.1 A sample stimulus adapted from figure 1 in Duncan (1984). Each object exemplifies two dimensions: Line (texture and tilt) and Box (size and gap location). In this stimulus, the line is solid and tilted right, while the box is large with a gap on the right.

The crucial result is the *same object advantage*: subjects were more accurate when the two features they reported were of the same object, rather than when they were split between two objects. If spatial attention is all that matters, then when both objects fall within the attentional spotlight, they should presumably be processed equivalently. Duncan's results suggest otherwise, for features that fall within the attentional spotlight are treated differently when they are features of the same object, rather than of different objects.³

Egly, Driver and Rafal (1994) developed another paradigm involving the presentation of two vertical or two horizontal rectangles on either side of a fixation point (Figure 7.2). A cue in the form of the outline of three sides of the end of one of the rectangles was flashed. After an interstimulus interval, a target was flashed in either a valid or an invalid condition, an adaptation of the spatial cueing paradigm: in the valid condition, the target is located near the location of the cue; in the invalid condition, the target is located at a distance from the cue. The subject's task was to report as quickly as possible the presence of the target, and, as expected, Egly et al. observed a reaction time advantage with valid versus invalid cues.

The novel result occurs in the invalid condition. Here, there are two cases: in the same-object case, the target appears opposite the cue, but in the



Figure 7.2 (A) Depiction of the Egly et al. object-based attention paradigm from Shomstein (2012). (B) Graph of the reaction times for the different conditions. Subjects showed the standard spatial cueing effect (first two bars) yielding a space-based effect (SBE) of attention. Subjects also exhibited an object based effect (OBE) where invalid cues within the same object showed faster reaction times. Figures reproduced courtesy of Sarah Shomstein and with permission of John Wiley and Sons.

same rectangle (the cued rectangle); in the different-object case, the target appears opposite the cue in the uncued rectangle. Crucially, in both cases, the target is equidistant from fixation and cue, so there should be no difference in regards of spatial attention. Relocating the spotlight in either invalid case would involve moving it the same distance from fixation or from the location of the cue. Egly et al., however, observed a same-object advantage relative to the different-object case. They describe this as the cost of having to switch attention from one object to another or, equivalently, the advantage of switching attention within an object. Similar results were observed if the object in question was seemingly partially occluded (Moore, Yantis, and Vaughan 1998). This paradigm is now a widely used method to probe the influence of objects on attention.

Psychologists speak of these effects on attention as object-based, and this raises a question: Do the results from Egly et al. demonstrate object attention, or do they show an influence of an object on spatial attention? After all, Egly et al. have essentially modified the spatial cueing paradigm. This might seem a subtle point, but the distinction can be clearly drawn. Imagine that spatial attention is fundamental and that so called "object" attention is really due to a modification of spatial attention. There is then a substantive difference between object attention and object-based (spatial) attention. Metaphorically, we can imagine two distinct mechanisms, a spotlight for spatial attention and a sticky finger for object attention, the former illuminating spatial regions, the latter latching on to objects (the finger metaphor is borrowed from Zenon Pylyshyn's work to be discussed in Section 7.4). If spatial attention is fundamental, then the Egly et al. effects are not evidence for the presence of an attentional finger, but only of a modification of the attentional spotlight. The effect is not due to the selection of the cued rectangle, with the finger extending to the object. After all, the task does not require that the rectangle be selected. Rather, cueing the rectangle activates an appropriate shape representation that alters the shape of the spotlight, as if one put a rectangular filter over it. Thus, objects affect spatial attention without themselves being the target of attention.⁴

Duncan's paradigm might be less susceptible to this worry since his task involved focusing on features of one or two objects, so selection of objects and their features is part of the task. There is another object attention paradigm that also seems less susceptible to this worry, namely the multiple object tracking (MOT) paradigm developed by Zenon Pylyshyn. As the rubric suggests, the task is directed at objects to visually track them. In this case, the empirical sufficient condition will be satisfied: subjects select objects for the purpose of tracking, so they are attending to those objects. Specifically, a subject observes a set of identical types of objects (circles or squares) on a computer screen (Figure 7.3). A subset of these is cued to identify them as the targets for tracking. The objects then move in an unpredictable way. When they come to a halt, subjects are asked to confirm that a given object was among the original targets or to identify all the targets.

Using this paradigm, Pylyshyn and coworkers showed that normal subjects are able to accurately track about four to five objects. Indeed, it seems plausible that attention is divided between distinct objects.

As Brian Scholl (2009) has noted, MOT is distinctive among the paradigms in requiring extended and active attention to a scene under conditions that



Figure 7.3 A depiction of the multiple object tracking paradigm, adapted from Scholl, 2001, p. 9. In the first panel, targets are identified by flashing them. They then move in a random fashion with subjects tasked with tracking the targets. After the targets come to a halt, subjects are asked to identify the targets, say, by moving a cursor to each.

echo real world tracking conditions. The paradigm can be manipulated in different ways yielding robust results. Pylyshyn, Scholl and others have established many of the parameters affecting performance in MOT. For example, subjects are capable of tracking targets even if the targets change in color and, within certain limits, change in shape while they move. Tracking can occur even if objects disappear momentarily behind an occluder (Keane and Pylyshyn 2006) or when the objects are spatially overlapping and do not move although their properties change over time (Blaser, Pylyshyn, and Holcombe 2000).

There are, however, limits to tracking. For example, when a line connects a target object with a distractor object, thus generating a more complex object with parts, tracking of the target is impaired even when each object moves independently (Scholl, Pylyshyn, and Feldman 2001). Scholl (2009) suggests that attention spreads throughout the complex object, making it difficult to track the target. Furthermore, when objects moved by breaking into smaller units and then reforming at a subsequent location (as if they were liquefied, then "poured" into the new location where they reformed), subjects were again impaired in tracking (VanMarle and Scholl 2003). This and other data suggest that MOT is object-based and is facilitated when targets behave like normal objects. Overall, MOT provides a clear paradigm where the target of attention is an object.

The next step is to connect object attention to demonstrative thought about objects. First, one needs to explain how attention makes the object available for thought. This focuses on a specific functional role of attention. Second, some have argued that conscious attention is necessary for demonstrative thought. Let us examine each step after examining how demonstrative thoughts are distinctive thoughts.

7.4 Demonstrative versus descriptive thought

Consider perception-based (or perceptual) demonstrative thoughts about an object, a thought entertained on the basis of a concurrent perceptual experience of the object and typically expressed by the use of demonstratives such as "this" or "that" ("demonstrative thought" henceforth means perception-based demonstrative thought).⁵ For example, imagine Jane pointing at a building in a skyline and saying, "That is the tallest of the lot." She uses the sentence to express what she is thinking, and a companion can understand her immediately by looking at what she points to. In contrast, imagine that in a different time and place, a stranger reads a written report of Jane's exact words. By merely reading the sentence, "That is the tallest of the lot," the stranger would not understand the specific thought Jane expressed. Why? A natural answer is that the stranger would in some sense not know which object Jane referred to with the demonstrative "that," something that Jane's companion knows by looking at what Jane refers to. This suggests two conditions for grasping Jane's perception-based demonstrative thought:

- (a) The know-which condition: one understands Jane's thought only if one knows to which building she referred;
- (b) The perception condition: seeing the building Jane was pointing to is necessary to understand the perception-based thought she expresses.

(a) is controversial, but I will accept it. Understanding a demonstrative thought, and yet not knowing which object the demonstrative refers to, seems to be incompatible. The central question is how to understand (b): Is seeing or some form of perception really necessary? Can't one express Jane's thought in the absence of seeing the building and meet the know-which requirement? After all, one can use a description: "the building pointed [or referred] to by Jane was the tallest of the lot." This allows one to refer to the same building and seemingly express Jane's original thought. The stranger who was not with Jane can use the same ploy.

Nevertheless, there are reasons to think that there are two different kinds of thought at issue: a demonstrative and a descriptive thought. To bring out the difference, let us assume that (a) whatever concepts are, they are the building blocks of thoughts, and that (b) the structure of sentences can reflect the conceptual structure of the thoughts they express.⁶ One might then think that the thought expressed with the descriptive sentence is conceptually complex, given the use of the definite description, "the building pointed to by Jane." In contrast, the thought expressed with the demonstrative seems conceptually simpler. After all, understanding the description requires having the concepts expressed by the words that make up the description ("building," "Jane," etc.), while understanding the demonstrative "that" in the context of use does not.⁷ One need only look at the building to understand what Jane says. Having the object in view in normal circumstances seems to be sufficient to grasp a demonstrative thought about the object. One can then initially distinguish between demonstrative versus descriptive thoughts: the former is conceptually simpler than the latter and depends on perception in a unique way. There is much more to say about the idea of descriptive versus demonstrative thoughts and much that is controversial. Our hypothesis is that there is a difference in the conceptual structure of these two forms of thought and that perception allows for grasp of the conceptually simpler demonstrative thought.⁸ Might attention also have an important role to play?

7.5 Is attention necessary for demonstrative thought?

We now consider the necessity of attention for demonstrative thought, first with general considerations regarding selection for action that suggest that attention is necessary, and then two proposals on more specific implementations that might explain how attention provides for demonstrative thought by playing a crucial role in making objects available in perception.

7.5.1 Thinking as action

In Chapter 3, I argued for a conception of attention as selection for action. From that perspective, the necessity of attention for perception-based demonstrative thought seems plausible. Note that demonstrative thought involves the deployment of concepts corresponding to the propositional content of the thought. Thus, a demonstrative concept and a predicate concept are joined to yield a complete demonstrative propositional content: that (ball) is kickable or that (building) is the tallest of the lot. Begin then with cases where one must figure out what to do or what to believe. In such cases, thinking is active, something done in the service of one's broader epistemic or practical goals. Here, there is a Many-Many Problem: so many possible items of thought and so many ways to think about those items. One's thinking will be constrained by current goals, whether determining what to kick or making judgments about skyscrapers. Since perceptionbased demonstrative thoughts are reliant on perception, yet demonstrative thoughts by their nature target only a subset of what is perceived, it looks like the specific dependence of a demonstrative thought on a specific aspect of perception will require input-thought coupling: the visual experience of a specific ball or a specific building is coupled to a thought formed on the basis of that experience. But this coupling, I have argued, entails attention, specifically attention to the target of the input state. Thus, while the previous discussion has emphasized a close tie between demonstrative thoughts and perception, it seems that it is specifically attention in perception which grounds demonstrative thoughts that are based on what is currently perceived. This highlights how the necessity claim falls quite naturally out of the selection for action view. Taking this claim under advisement, let us then consider two more concrete suggestions about possible ways that perception grounds demonstrative thought through attention.

7.5.2 Fingers of instantiation (FINST)

Zenon Pylyshyn has taken MOT to highlight a fundamental and basic visual capacity for individuating objects. Indeed, Pylyshyn notes that MOT was developed to test a conceptual assumption, namely, that behavior and thought depend on a primitive capacity for object individuation. Pylyshyn invoked a metaphor of sticky, flexible visual fingers that could glom onto and track objects in the world. If such fingers existed, what he called fingers of instantiation or FINSTs, they should be deployed in visual tracking. So, FINSTs were postulated on roughly *a* priori grounds and empirically tested with development of the MOT paradigm. Tracking in MOT deploys the use of multiple FINST tokens, one for each target tracked.

Pylyshyn (2007) argues that FINSTs serve as a basic mechanism of reference, something that explains representation, but is not itself representational. Specifically, FINSTs are a mechanism by which the visual system locks onto objects in a way that is independent of representing the locations and properties of those objects. On Pylyshyn's account, FINSTs provide a connection between the world and visual object representation by providing a link to object files that can store information about the objects that are the targets of FINSTs. It is in the object file that representations of an object and its properties first gets a grip, and Pylyshyn speaks of the contents of the object file as conceptual, hence representational.⁹ FINSTs provide a selective informational channel linking a specific object to a conceptual representation that can then be deployed in thought.

Earlier, I distinguished between the causes of attention and the targets of attention. Pylyshyn invokes a similar distinction between the triggering of a specific FINST token and the encoding in the object file of the property that triggers the FINST. The crucial point is that a FINST can be triggered by properties of its target without necessarily leading to the encoding of those properties in object files. In light of experiments that show that subjects do not retain much information concerning the properties of objects during MOT, Pylyshyn notes that "nothing is stored in the object files under typical MOT conditions" (Pylyshyn 2007, 40), though information regarding those properties could in principle be subsequently stored.¹⁰

The facts that FINSTs serve object tracking and that object tracking involves attention suggest that FINSTs are tied to attention. Brian Scholl (2009) notes that "there may be nothing to MOT beyond attention" (55), and this suggests that we should think of FINSTs as part of the mechanism of attention. In contrast, Pylyshyn treats FINSTs as a preattentive mechanism, something that precedes attention. For Pylyshyn, visual attention is directed at visual objects, and this requires visual object representation. Given that FINSTs serve the role of allowing for the representation of objects (via an object file), they provide the materials on which attention operates. So FINSTs are earlier in the visual processing hierarchy than attention, and the deployment of a FINST does not constitute the deployment of attention. Resolving the issue whether FINSTs are preattentive or attention goes beyond our purview. Let us just note that if Scholl is right, then we have a proposal as to why attention is necessary for demonstrative thought, namely, that attention makes object representation possible (cf. Treisman's Feature Integration Theory where attention is the glue that binds features together for object representation, Chapter 1). Without attention, we would not be able to individuate objects so as to think about them.

7.5.3 Selection and access; figure and ground

John Campbell (2011) has also made proposals on how attention makes objects available for thought. His account emphasizes a distinction between two aspects of attention: *selection* versus *access*. Recall that the gatekeeper view asserts that one is conscious of a property only if one can or does access the property. Campbell, however, suggests that consciousness should be tied to selection. He argues that there are cases where we have selection of a property in perception even though the subject is incapable of accessing the property. Nevertheless, the subject is conscious of the property, so consciousness outstrips access. Campbell suggests that this is true of young children with respect to color. While their behavior demonstrates that they are aware of color, children pass through a stage where they cannot reliably apply color terms. They are not able to access color for verbal report and, plausibly, for other tasks that require an ability to conceptualize color. Campbell suggests that this is a case where a property can be selected in vision without being accessed by conceptual processes. On this account,



Figure 7.4 Black and white photos of Ishihara diagrams used to test for color blindness. In the right diagram, a faint 12 can be discerned. No number can be seen in the left diagram (the number 74 would be visible in the original colored diagram). Subjects normally use color to individuate the object from background. In these two cases, one is blind to the number on the left (what I called individuation blindness in Chapter 5), and one can barely select the number on the right. Photographs kindly provided by Cameron Wittig and used with his permission. © Cameron Wittig.

there is a selective function of attention that is for consciousness, even if it does not (yet) serve action.

Campbell argues that attentional selection of color allows the subject to individuate objects in experience, specifically by distinguishing objects as figure from background (figure/ground segregation). In this way, attention can provide objects for thought. Consider Ishihara number-dot displays that are used to test for color blindness where a numerical figure is made up of dots in one color, the background in another (Figure 7.4).

For subjects with normal color vision, the number can be clearly distinguished from the background on the basis of a difference in color. That is, the selection of the color of objects allows the subject to individuate that object from background. For those who are colorblind, in that they are not able to distinguish two distinct colors, the number is not seen because it cannot be separated from background. Colorblind individuals are not able to select colors in the relevant sense. In discussing such displays, Campbell notes that "if you see the thing at all, you had to be using color to select it" (330). Awareness of objects then requires attentional selection of its properties.

There are two ways to understand the appeal to selection vis-à-vis figureground segregation. On the one hand, a natural idea is that the relevant figure-ground segregation is a basic visual computation that is needed if

one is to have the visual representation of an object. At some point, that object must be visually individuated from other objects and from the background. On the other hand, figure-ground segregation might refer to the phenomenology of seeing a figure as against a background, and in that way individuating it from the background in one's experience. I shall return to the second, phenomenal conception of figure-ground segregation in a later section, but let us briefly make an observation concerning the mechanisms of selection that serve figure-ground computation as the relevant proposal for attention. On this reading, attention serves as a basic mechanism for visual processing that leads to the individuation of objects by vision. In that way, it plays a similar functional role as a FINST in that it counts as a causally necessary condition for the individuation of objects if there is to be visual representation of objects. The question then is whether the process of selection needed for figure-ground segregation counts as attention or simply as a selective mechanism. It is worth noting that the connection between figure-ground segregation and attention has long been debated in psychology. This is a complicated issue that we cannot delve into here, but a central question is whether attention does influence figure-ground segregation or whether such segregation is independent of attention (for a recent overview, see Kimchi 2009; for a recent experiment that suggests that segregation can be independent of attention, see Kimchi and Peterson 2008). Moreover, we might wonder whether figure-ground segregation in this sense isn't really just a subpersonal process that is necessary for personal level awareness of objects.

There are many empirical issues that arise when considering FINSTs and figure-ground segregation as two sorts of visual processing that are connected to visual object representation. The interesting question is how they are related to attention, and opinions differ. Perhaps the most ecumenical approach here is to fall back on the selection for action account as the broadest conception of attention. In doing so, we can see FINSTs and figure-ground segregation as two mechanisms that contribute to the needed object-thought coupling if we are to have demonstrative thought. That is, those processes would arise in a discussion of how selection for action is implemented in creatures like us. In this way, the two mechanisms are tied to attention. There is more work to be done here in theories of the specific role attention plays in demonstrative thought that are enriched by appeal to fundamental mechanisms of object individuation. These are issues that remain open, subject to collaborative work in cognitive science and philosophy. The central point for our discussion is that the three proposals discussed in this section point to attention as necessary for demonstrative thought via making objects available for thought. When attention is sufficient for demonstrative thought is a key issue to which we now turn.

7.6 When is attention sufficient for demonstrative thought?

Thinking generates thoughts in a way that is constrained and motivated by one's intentions. To think about a specific item perceived, attention must couple thought to that very thing. In this case, attention serves to ground an information link between thought and the world so as to make possible a special way of thinking about the world. This section considers cases that point to a critical role for *conscious* attention and then examines three recent philosophical theories of demonstrative thought that provide a fuller story of how conscious attention yields demonstrative thought. The issue to be addressed is what more must be added specifically to attention so that it can be part of a sufficient condition for demonstrative thought, whatever further non-attention conditions are required (e.g., that the subject have certain conceptual capacities, that normal conditions obtain, etc.). Mere attentional selectivity of the unconscious variety might not be sufficient for attention to play its necessary role. Attention might need to be conscious.

7.6.1 Blindsight and the sea of faces

Can unconscious attention serve to fix demonstrative thought? Recall that attention-mediated coupling can be unconscious, as in the case of the blindsight patient GY (see Chapter 4, Section 2). Yet, when GY offers guesses regarding objects in his blind field, selecting those objects to inform his reports, is he having the demonstrative thought that that object is oriented vertically? This seems unlikely, at least on its face. He can't see the object in the way that we do. More likely, his thought is a descriptive one: the object (probably!) in my blind field is oriented vertically. In tapping into accurate information about that object, I suggested that GY is unconsciously attending to those objects, i.e., selecting the object in unconscious vision. Yet despite this, some feel that the deployment of attention by GY does not suffice for a demonstrative thought, even if it supports a descriptive guess. What is missing?

Perhaps what is missing is confidence and spontaneity. Imagine a blindsight patient growing more confident as he understands his condition, recognizing

that his blindsight is a reliable source of information about objects in his blind field. Given this confidence, as he navigates the world, he might just start producing "guesses" voluntarily, reports that he expresses with verbal demonstratives. "That is red," (is vertically oriented, square, etc.), he says. But despite his use of a demonstrative in his utterances, does he express a demonstrative thought? Why not say instead that he gives verbal expression to the more complex descriptive thought attributed to GY? Our cognizant blindsighter has learned to use the word "that" appropriately, but the thoughts he is capable of entertaining remain the more complex descriptive thoughts. His confidence and his appropriate use of the word "that" spontaneously and accurately does not suffice to allow him to move from a descriptive thought to a demonstrative thought. The object remains, in some sense, perceptually at a distance from him. If you have this intuition, then even if the cognizant blindsighter deploys attention to make accurate judgments of objects he does not consciously see, he continues to only have descriptive thoughts about attended objects. The obvious bit that seems missing from GY and the cognizant blindsighter is consciousness.

On the other hand, some might have the opposite intuition: the cognizant blindsighter can entertain demonstrative thoughts about objects in his blindfield. Imagine the cognizant blindsighter moving around the world, and making judgments using demonstrative expressions. Imagine that when queried which object he means to refer to, the blindsighter points at a specific moving object, tracking it with his finger. He admits that in the old days before cognizance, he would have felt like he was guessing, but now, he understands his condition and simply refers and points. If this subject can keep tabs on the object via unconscious attention in a way that generates spontaneous and reliable judgments and actions, is it not plausible that the attention-mediated coupling does allow for a demonstrative and not merely descriptive thought about the object?¹¹ Although I will discuss demonstrative thoughts under the assumption that the blindsighter is not capable of entertaining them, one should keep the alternative firmly in mind.

Consciousness seems to some to be a crucial ingredient in making demonstrative thought possible. So, does the introduction of consciousness back into the mix secure demonstrative thought? Even this might not be sufficient. John Campbell raises a relevant challenge in his *sea* of *faces* case. Imagine being at a crowded dinner party where a companion singles out a woman with the expression "that woman". Normally, to understand the demonstrative, one visually locates the woman in question, and in that way

224 ATTENTION AND DEMONSTRATIVE THOUGHT

knows to whom the companion refers. Campbell asks us, however, to imagine that our visual experience is as of a sea of faces such as when looking at a crowded stadium. One has not yet visually picked out the person in question. Now imagine that like the blindsighter, one is able to offer reliable guesses when queried about this person. This comes as a surprise. The judgments seem to be just lucky. In this case, one might feel unable to grasp the demonstrative, "that woman," despite conscious vision and unconscious attention. Campbell writes:

It is only when I have finally managed to single out the woman in my experience of the room, when it ceases to be a sea of faces and in my experience I focus on that person, that I would ordinarily be said to know who was being referred to. So it does seem to be compelling to common sense that conscious attention to the object is needed for understanding of the demonstrative.

(9)

The crucial element then is conscious attention, i.e., attention with a specific phenomenal upshot. Campbell often refers to the highlighting of the object by attention. This means that the story of attention's role in fixing demonstrative thought is more complicated than simply noting that attention selectively allows objects to be coupled to thought. For some philosophers, attention must be conscious. We now explore the recent work of three philosophers who provide accounts of what is necessary for demonstrative thought beyond the attention-mediated coupling of perception to thought. All three accounts emphasize conscious attention, but they are also distinctive in that they invoke a *normative* dimension. While normativity is a philosophically loaded concept, it is used here in a fairly thin sense to capture a common feature of these accounts: namely, their appeal to some notion of justification.

7.6.2 Attention and knowledge of reference: Campbell

John Campbell (2002) provided an important systematic treatment of how attention grounds demonstrative thought.¹² For Campbell, the relevant notion of attention is not merely an information-processing notion, but also a phenomenal one. Conscious attention of an object involves "experiential highlighting of the object" (p. 2). In visually attending to an object, a subject separates it from the background, rendering the attended object as

in the foreground (p. 25). This suggests that there is a distinctive phenomenal character tied to conscious attention, or at least a phenomenal feature of those forms of attention that can serve demonstrative thought. Attention phenomenally individuates an object so that one can think demonstratively about it.

Campbell's account of how attention provides for demonstrative thought appeals to a connection between attention and knowledge of meaning and reference. Recall the know-which requirement: one understands a demonstrative only if one knows which object the demonstrative refers to. Campbell's proposal is that it is conscious attention that provides for, indeed constitutes, knowledge of which object is being referred to. To argue for this, Campbell begins with what he calls the Classical View that provides an account of the cognitive role of one's knowledge of the meaning of a term. We can think of such knowledge as guiding one's response to a proposition. Specifically, "knowledge of what it is for a proposition to be true is what causes, and justifies, your use of particular ways of verifying, and finding the implications of, that proposition" (24). For example, one's understanding of a scientific hypothesis (proposition) both guides and justifies how one goes about experimentally testing it or the theoretical consequences one draws from it. Thus, certain experimental procedures a scientist undertakes make sense given what the hypothesis claims. Correspondingly, a student who does not understand the proposition lacks knowledge that guides appropriate behavior. Applied to demonstratives, the Classical View holds that knowledge of the referent of the demonstrative is what causes and justifies how one goes about verifying or responding to demonstrative propositions. So, knowing which building Jane is pointing to guides and justifies how one goes about testing whether it is really the tallest of the lot. For example, one might compare other visible buildings in relation to it or take measurements from it. Doing so makes sense, given the thought Jane expresses.

Knowledge of demonstrative reference plays both a causal and justificatory role, and Campbell argues that attention fills this role for demonstrative thought. The causal dimension is captured in his causal hypothesis:

When, on the basis of vision, you answer the question, "is that thing *F*?", what causes the selection of the relevant information to control your verbal response as well as the maintenance of this link is your conscious attention to the thing referred to.

(13)

On this view, conscious attention, a personal-level phenomenon, affects subpersonal processing to serve a specific task, namely answering a question. Attention does so by setting the goal of information processing, which thereby brings it in line to serve a task. It identifies the target towards which such processing is directed. Critically, it is only conscious attention that is capable of target setting. For example, in manipulating an object, conscious attention identifies the target to which motor processing is directed in order to generate movement towards that object.¹³ The justificatory dimension is then cashed out through the goal of information processing as determined by the object of attention. In selecting an object, attention thereby sets the standard for assessing the success of information processing. So, if an attended object is the target of a reaching movement, the accuracy of reaching is assessed in respect of that target: In light of the target, is the resulting trajectory correct? Is the prepared grip appropriate to the structure of the object?

One way Campbell fleshes out the proposal is in terms of Anne Treisman's Feature Integration Theory (FIT; Campbell is not tied to FIT, and in more recent work he appeals to a Boolean Map Theory of attention (Huang and Pashler 2007)).¹⁴ According to FIT, the visual system encodes different features in visual space via distinct feature maps, and attention can bind those features together to yield an object representation. Campbell notes two uses of feature maps where attention is involved. First, a "low-level," not necessarily conscious form of attention aids in the binding of features to allow for object representations (recall his distinction between selection and access). Second, once an object representation is secured, conscious attention to the object allows the subject to tap into information available in feature maps for use in verifying propositions about the object in question, say, whether it is red or square. The story of attention, then, is that it sets a target for processing that allows for the second use of features maps in verifying demonstrative propositions about that target, a use that is justified or appropriate given that attention targets the object in question.

When Jane points to the building and says, "That is the tallest of the lot", one can verify her proposition by consciously attending to the building. Attention triggers a causal process that binds relevant features of the building together to generate a representation of the building and makes available information about relevant features from feature maps to inform one's subsequent judgment. It is then appropriate (justified) to draw on feature maps in this way, given that the target of processing is the target of attention: it is the building to which Jane demonstrates. That building is

the goal of information processing and thus the standard for assessment of subsequent activity. Accordingly, attention to an X can both cause and justify the ways that one verifies propositions about X.¹⁵ Since attention plays the role of knowledge of reference, Campbell takes it as constituting knowledge of reference of the demonstrative. The necessary role of attention then is not just in its causal dimension, its yielding a coupling of perception to thought, but its doing so in a way that sets a standard for information processing that serves our practical and epistemic goals. To do this, attention must be conscious, and for this reason, the cases of standard blindsight, cognizant blindsight, and the sea of faces do not identify a form of attention that allows for demonstrative thought, since they do not involve conscious attention.

7.6.3 The rational role of conscious attention: Smithies

Declan Smithies (2011a) has also argued that conscious attention is necessary for demonstrative thought. Like Campbell, he sees attention's role in this context as having a normative dimension. Yet while Campbell is concerned with attention setting standards for information processing in a way that is tied to knowledge of reference, Smithies emphasizes a link between attention and justification. He argues for the necessity of attention in demonstrative thought as follows:

- One has a demonstrative concept of an object O only if one has information about O which provides immediate, defeasible justification to form beliefs about O, and which one is able to use in forming immediately justified beliefs about O;
- (2) One has information about O which provides immediate, defeasible justification to form beliefs about O, and which one is able to use in forming justified beliefs about O, only if one has conscious perceptual attention to O;
- (3) One has a demonstrative concept of an object 0 only if one has conscious perceptual attention to 0.

(Smithies 2011a, 32)

The concept bridging the premises is the notion of immediate justification for beliefs. Central to Smithies's conception of justification is a necessary role for consciousness in explaining how perception justifies belief. For current purposes, we can think of perceptual justification in terms of perception's providing reasons for beliefs, where a reason is a consideration that speaks in favor of a belief. Thus, one's perceptual experience of a thunderstorm can provide reasons for believing that there is a thunderstorm. Given this conception of perceptual justification, there are two relevant questions. First, why must perception be conscious if it is to play a justificatory role? Second, why must attention be conscious if it is to play an epistemic role in perceptual justification? For Smithies, each question is tied to two separate senses of justification: (1) perception as having the property of providing reasons for belief (propositional justification); and (2) the use of propositional justification in forming a belief (doxastic justification). Let us briefly consider propositional justification.

Again, consider blindsight. As noted previously, blindsighters can make accurate reports of their environment on the basis of vision, at least when prompted. Still, blindsighters seem to be in a deficient epistemic position regarding the world relative to normally sighted individuals, and one basis of this deficiency seems to be the absence of visual consciousness. Smithies elaborates the grounds for the intuition by noting that blindsight states do not provide for immediate justification for beliefs, while normal conscious visual states do. Perceptual justification is "immediate in the sense that it does not depend on my having independent justification for background beliefs about the reliability of my experience" (Smithies 2011a, 23). When I have a good view of a ball and conditions are otherwise normal, I am perceptually justified in forming beliefs about its shape, namely, that it is round, without the need to justify the reliability of my experience in these conditions. A blindsighter confronted with the same ball in the same conditions can accurately report that the ball is round when prompted, though this report will feel like a guess. The fact that the report feels like a guess suggests an absence of available justification. To the blindsighter, the report feels like a leap in the dark, not something that is grounded in a reason.

Are blindsight subjects guessing? Consider three notions of guessing. First, blindsighters have the phenomenology of guessing, for it feels to them as if their judgments are shots in the dark. They are keenly aware of the difference between the judgments that they issue on the basis of normal perception of objects in their intact visual field and the guesses they are forced to issue in respect of objects in their "blind field". They feel at a loss with respect to their blind field. This leads to a second notion of guessing, for even if the blindsight subject issues a correct judgment about an object in the blind field, the blindsighter will not be able to articulate the reasons for the judgment. The idea here is that the absence of phenomenal consciousness of the object undercuts access to a sort of justification that is present in the normal case. In that respect, felt guessing corresponds to genuine guessing. The judgment is not based on reasons accessible to the subject. It is this case that is central to Smithies' discussion. Nevertheless, and this is the final notion, the blindsighter is not rendering a judgment on the basis of no information. After all, the blindsight judgment is informed by relevant visual information. The judgment is not like the guess of those who have closed their eyes or of the congentially blind who cannot in any sense see. Guesses made without access to reasons, and guesses made on the basis of no information, highlight a question about the type of access to an object required for demonstrative thought: Must such access be conscious or can it be based on unconscious information links?

Blindsighters can improve their epistemic standing with the help of further beliefs. Cognizant blindsighters might be aware of their condition and recognize that blindsight states are reliable sources of information about the environment. Accordingly, cognizant blindsighters are justified in forming the belief that the ball is round on the basis of their blindsight states. They know that their blindsight is a reliable source of information, and they can tap into it to form beliefs. For them, their reports are not guesses. Still, there remains a salient difference between cognizant blindsighters and normal perceivers, for the justificatory role of perception in the former is mediated, dependent on beliefs about the reliability of blindsight, whereas for normal perceivers, perception can provide immediate justification for belief. Reflection on blindsight suggests that it is consciousness that explains the capacity for immediate justification provided for by perception. Leaving matters at this intuitive level, let us ask: Why is conscious attention necessary in order to use the justification that conscious perception makes available?

Smithies notes that there is a distinctive phenomenology associated with attention: "there is a phenomenal distinction to be drawn between the attended foreground and the unattended background of one's experience" (Smithies, op. cit., p. 28). The question then is whether the distinctive phenomenology of attention implies a distinctive epistemic role. Smithies' description recalls Campbell's talk of experiential highlighting, and as we noted earlier, Campbell has argued that attentional selection is necessary for separating an object from its background. So, one might think that the

epistemic role of attention is to bring objects to the fore in experience to enable epistemic access to them. This capacity is what differentiates us from the blindsighter and the subject in the sea of faces case. Smithies' emphasis is on attention as allowing the subject to use perception to justify belief (doxastic justification). As he puts it:

attention plays a role not only in modulating the contents of experience, but also in formatting the contents of experience in such a way as to make them accessible for use in conceptual thought. Thus, attention to an object is necessary for converting the contents of experience into the contents of justified belief.

(30)

Smithies later elaborates the idea of formatting the contents of experience in terms of solving the "binding problem" at the level of conceptual thought, namely, by making possible the formation of beliefs where certain properties are predicated of the attended object. The proposal is suggestive. It is clear that Smithies focuses on attention as a form of access to the propositional justification perception provides. In any event, Smithies' discussion provides the basis for premise (2) of his argument, namely, that immediate justification implies conscious attention. Premise (1) then is a claim about the possession conditions for demonstrative concepts. Possession conditions for a concept specify what conditions subjects must satisfy if they are to count as having the concept. The issues concerning concept possession are also outside our purview, so let us simply summarize Smithies' premise: that possession of demonstrative concepts implies the availability of immediate justification. Given (1) and (2), it follows that to have a demonstrative concept of an object O, one must consciously attend to O.

7.6.4. Reference, justification, and attention: Dickie

Imogen Dickie (Dickie 2011) has also presented an intricate and distinctive account of the role of attention in demonstrative thought.¹⁶ Let us highlight two elements in her account that also emphasize the causal and normative dimensions of attention's role. On the causal dimension, Dickie emphasizes that attention must be what she calls luck-eliminating: where demonstrative reference succeeds in referring to an object, *O*, this is not merely a matter of luck, and where it fails in referring to *O*, this is a matter

of being unlucky. By way of an example of a luck-eliminating process, think of any skilled activity as deploying capacities that are luck-eliminating, for when an agent acts skillfully, the success of that act is not fortuitous. The deployment of a skill, even basic skills like reaching for an object, sets in motion a reliable process such that when the skill is successfully executed, this is not a matter of luck, and when the process goes awry, this is often a matter of being unlucky. The first idea then is that attention is this sort of luck-eliminating process in the case of demonstrative reference, essentially a skill that enables reference to objects. Consider in this context Pylyshyn's discussion of FINSTs as a mechanism that links the world to object files, where the latter are to be understood as conceptual representations. Dickie draws on this work, conjecturing that perceptual demonstrative concepts are object files of a certain provenance, namely, those whose informational content is provided for by attention. The reliability of tracking as experimental work shows seems to correlate with conditions where the target of tracking is an ordinary object: MOT targets move through a spatiotemporal path in a predictable way, with coherent changes to their structure. It seems plausible to conclude that attentional tracking is reliable when and only when the target of tracking is an ordinary object.¹⁷ The experimental data suggests that the ability to attend to an X is precisely luck-eliminating when that X is an ordinary object. Attention is just the sort of capacity to deploy when one wants to refer to an ordinary object.

Dickie also provides a general argument for normative conditions on reference and intentionality. Her argument for this is complicated, so I shall give the intuitive version, drawing on two ideas. The first connects intentionality to truth, and the second connects truth to justification. On the first, there are clear ways that intentionality is connected to truth, as when we give accuracy conditions as a way to spell out representational content: the belief that O is F is true if and only if O is F. Thus, the truth of the belief is tied to what the belief is about. On the second, truth seems to be connected to justification in that the better the justification for our beliefs, the more they are likely to be true. Thus, by (1) connecting intentionality to truth, and (2) connecting truth to justification, intentionality is connected to justification. In the specific case of demonstrative thought, a condition on demonstrative reference is justification that is conducive to successful reference.

Dickie's analysis of justification connects it to intention. Thus, while Smithies focuses on theoretical justification (justification of belief), Dickie focuses on practical justification (justification given intention). Further,

while Campbell sees attention as setting the standards for information processing, for Dickie it is intention that sets the standards that constrain attention. Consider Robin Hood who is a skilled archer and intends to hit a wand stuck into the ground at some distance. Given his intention, he deploys his archery skills and hits the wand. The justification for his deployment of the specific skills that he uses is due to the standard that his intention sets, namely, the act he aims to perform. Given this standard, his deployment of his archery skills is justified, since those skills are luckeliminating with respect to shooting an arrow at the wand. The skills are a reliable route to his goal: if he hits the wand, he is not lucky, and if he misses, then he is unlucky (e.g., a strong gust of wind blows the arrow off course). The point is that the relevant skill is not just causally efficacious in fulfilling one's goals, but that those skills are what one ought to deploy given one's goals. In general, specific practical goals justify the deployment of specific skills. The same structure holds for demonstrative reference. Consider Campbell's sea of faces case or Jane's talk of a specific building. To understand what the speaker is saying, what is being referred to, one must locate a particular target. Communicative intentions often have reference as a goal, to locate and think about a given object. In both cases, it is appropriate to deploy attention to fix reference, for it is a reliable process that locks on to an object. Given referential goals, the deployment of attention to a specific target is justified. Attention is a reliable, luck-eliminating process that allows the mind to lock on to ordinary objects, say, in attention-mediated object files.18

The issues raised in this section are complicated and challenging, but let us draw some lessons about the role of attention in demonstrative thought. Campbell, Smithies, and Dickie agree that an information link is not sufficient for demonstrative thought. In providing an account for how attention can fix demonstrative thought, each focuses on a functional role for attention that ties it to different notions of justification, whether setting standards for judging the appropriateness of information processing (Campbell), making available the use of propositional justification (Smithies), or being connected to the subject's referential intentions that justify the deployment of certain luck-eliminating (indeed referential) skills (Dickie). These accounts address what attention must ultimately bring to the table if demonstrative thought is to be established. A normative dimension is one additional element in the picture; a second is the role of conscious attention. Let us conclude with the question of conscious attention.

7.7 Is conscious attention necessary for demonstrative thought?

In arguing for the necessity of attention for demonstrative thought, I deployed the selection for action account. That account, however, allows that attention can be unconscious, and this raised the question whether the cognizant blindsighter can have demonstrative thoughts. For Campbell, Smithies, and Dickie, it seems that cognizant blindsighters cannot, for, minimally, they lack conscious attention. But is conscious attention really necessary for demonstrative thought? To set the final issue for this chapter, let us consider the following comment by John Campbell:

If I am to understand a demonstrative referring to an object, it is not enough merely that the object be there somewhere in my visual field; I have to attend to it. But the attention that is needed here is, as it were, a matter of experiential highlighting of the object; it is not enough that there be some shifts in the architecture of my information-processing machinery, remote from consciousness.

(Campbell 2002, 2)

Campbell's talk of experientially highlighting an object recalls the phenomenal conception of attention discussed in Chapter 4. That is, attention is conscious in the sense that there is a what-it-is-like to attend to an object, a distinctive phenomenology of attention, and it is conscious attention that is needed for demonstrative thought.

In discussing the phenomenal conception of attention, I suggested a different interpretation of James's talk of attention as the focalization and concentration of consciousness, namely, that attention is a way of responding to consciousness and not itself a mode of consciousness. This suggests two proposals for the role of attention in demonstrative thought:

- (a) Attention is necessary for demonstrative thought, but it need not be phenomenally conscious;
- (b) Specifically conscious attention is necessary for demonstrative thought.

For (a), the additional element is just attention in the sense of coupling conscious perception to thought. For (b), attention must also have a phenomenal upshot. Let us grant for the moment that conscious perception is necessary for perception-based demonstrative thought. The question then is

this: Once we have allowed that, why is an additional phenomenal element needed, namely, some attentional phenomenology? Why wouldn't conscious perception, plus relevant attentional coupling between an object and thought, be all that is needed with respect to attention?

Once we concede that consciousness in perception is necessary, an argument is needed for the additional phenomenal element that attention is supposed to bring. Smithies provides an explicit argument that attention is a distinctive mode of phenomenal consciousness, and this suggests something like a corresponding attentional phenomenology:

- (1) Attention is what makes information fully accessible for use in the rational control of thought and action;
- (2) But what makes information fully accessible for use in the rational control of thought and action is a distinctive mode of consciousness;
- (3) Therefore, attention is a distinctive mode of consciousness (Smithies 2011b, 248).

This argument leads to a distinctive conception of attention as rational access consciousness. Specifically, Smithies distinguishes two types of functional role that attention might play: a broader causal functional role where attention causally influences behaviour, and a narrower rational functional role where attention influences rational processes such as the justification of belief. A theorist who takes attention as selection for action endorses a broader functional role for attention and allows for unconscious attention as capable of controlling behavior. This theorist takes Smithies' rational control as a special case of selection for action. Smithies' narrower position holds that attention is associated with rational role, and any process that plays only the causal role is merely an ersatz form of attention. So, the selective processes we noted in discussing the blindsight patient GY only demonstrate ersatz attention, not genuine attention.

But why load the explanation in this way? Why not say that consciousness is all in perception, and not in attention. One might question Smithies' argument by invoking the distinction between accessibility and access: attention isn't what makes conscious experience accessible for thought, consciousness in perception is sufficient to do that; attention in the context of thought entails thought's access to the deliverances of visual experience. For in attending to an object to fix a thought, the object is coupled to thought and, in that sense, accessed by thought. Experience's being phenomenally conscious is sufficient for the experience being accessible to thought (contrast blindsight); one only then needs to attend in experience to access what is consciously perceived. Thus, one might replace Smithies first premise with the following:

(1A) Attention is what makes information fully accessed in the rational control of thought and action.

What makes information accessible, as per (2), is conscious perception. In this case, the argument will not go through. The central disagreement then is as follows: for Smithies, accessibility is tied to attention; for his opponent, accessibility is tied to conscious perception, while access is tied to attention.¹⁹ How might we adjudicate this disagreement?

On Smithies' account, what is attention doing to alter consciousness such that it is now accessible to thought? Like Campbell, Smithies also speaks of attention phenomenally rendering an object as figure to ground. The contrast here is Campbell's sea of faces example, for before attention is deployed, objects are not individuated from ground. A sea of faces does not allow for a single face to be accessible; rather, we have to individuate the object from its background by attending to it. Only then is the object accessible. How compelling is the sea of faces case in establishing the need for conscious attention? I am severely myopic. When I take my glasses off, the faces in a photograph in front of me become blurred, a sea of faces. Yet I can also lock my eyes onto a part of the sea, say by selecting a property, perhaps a patch of color. The faces still do not resolve, but I think about that thing (we can allow, to eliminate guesswork here, that I know that the color patch I am attending to corresponds to the color of a specific face). The idea here is that certain ways of attending to a feature suffice for attending to the object with that feature. Or do the exercise in a different way: imagine focusing on an object so as to think about that object, and then, perhaps through manipulation of the air or of your brain, blur things so that objects become indistinct. In focusing on a property that remains visible, one might still maintain a lock on that object. So, the proposed phenomenology might be decoupled from, and thus not necessary for, demonstrative thought. What remains is an information link that attention establishes, connecting conscious perception to thought.²⁰ There are many questions to be answered here, but let me emphasize the minimal position which should be the default position since all parties in the debate can agree to it: vision-based demonstratives require conscious vision and attention suitably deployed. The remaining question is whether attention so deployed must itself be conscious.

7.8 Summary

This chapter has investigated object attention and how attention to objects serves perception-based demonstrative thoughts about objects. Attention might do so by making objects available in perception, perhaps by affecting basic perceptual processing, or it might do so as part of a general condition on thinking as an activity where attention mediated coupling links an object to one's thinking. Still, while many acknowledge that attention is necessary for perception-based demonstrative thought, there remain several questions that still need addressing:

- 1. While attention is necessary for demonstrative thought, how precisely does it provide for demonstrative thought (i.e., what is a sufficient condition for demonstrative thought)?
- 2. Is consciousness in fact necessary for demonstrative thought? That is, might our cognizant blindsighter be capable of demonstrative thought?
- 3. Is conscious attention necessary for demonstrative thought? Might nonphenomenal attention plus conscious perception be enough?

The three philosophers we have discussed have made the first foray in exploring the central role of attention in grounding demonstrative thoughts about objects. This provides a firm foundation to make an attack on these essential questions about our ability to think about the world.

Suggested reading

Evans (1982) is an early, difficult, and richly rewarding discussion of demonstrative thought. Campbell (2002) provides an important discussion of reference where attention is front and center. Dickie (2011) and Smithies (2011a) provide recent philosophical theories of demonstrative thought where attention plays a central role, and both have forthcoming monographs discussing the issue. Pylyshyn (2007) presents his theory of fingers of instantiation (FINSTs) and applies it to explaining thought about the world. Scholl (2001) provides a helpful discussion of object attention and Scholl (2009) examines multiple object tracking (MOT) with emphasis on attention. Shomstein (2012) provides an overview of the possible mechanisms underwriting object-based attention.

Notes

- 1 Another target we shall not discuss is attention to time. There is some recent work by Ian Phillips (2012) with reference to interesting psychological experiments on this issue.
- 2 There are other *dependency* conditionals that might also seem plausible at first glance.

If S attends to object O which is located at location L, then S attends to L.

If S attends to feature F of O, then S attends to O.

If S attends to empty space L, then there must be some feature F at L that S attends to.

- 3 See also Baylis and Driver (1993) and Baylis (1994). Ulric Neisser and Robert Becklen (1975) made an early contribution in identifying non-spatial forms of attention. Neisser and Becklen aimed to provide a visual analog of some dichotic listening experiments. Specifically, they provided subjects with two streams of visual information to the eyes as auditory attention theorists provided two verbal streams to a single ear. Neisser and Becklen superimposed two videos, one video of two pairs of hands playing a slapping game (player A lays hands on player B's hands where B tries to slap A's hands, including making feints as well as actual attempts), the other video of three individuals passing a basketball back and forth. Both videos fall within the same spotlight of attention. Subjects either had to count the number of slapping attempts or the number of passes (this paradigm was later adapted in inattentional blindness work, discussed in Chapter 5). What Neisser and Becklen showed was that subjects could selectively track the relevant events with very high accuracy. This accuracy was drastically reduced, however, when subjects were required to attend to both events. Again, there seems to be a difference in response within the same spatial region. Subjects could selectively attend to events, such as the passing of a ball or the slapping of a hand, even when these events were superimposed spatially. Similar results were demonstrated in static displays regarding overlapping colored figures by Irvin Rock and Daniel Guttman (1981).
- 4 Two models of object-based effects have been discussed recently. On the first model, the *attentional spreading hypothesis*, once an object is cued, attention automatically spreads within the confines of the cued object. Sarah Shomstein (2012) has suggested that this automaticity has two

components: the construction of a spatial gradient of attention at the point of the cue, and the top-down modulation of the spread of attention within the object (Martínez et al. 2006, 298), with the object representation being a determinant of top-down influence. An alternative model, the *attentional prioritization hypothesis*, endorses the first form of automaticity, namely the automatic spread of attention from the cue, but argues for a more flexible deployment of attention in respect of the object. Thus, while attention can spread selectively within a cued object, attention can be variably deployed given task demands. We can think of this as a form of top-down modulation that is goal-driven and, hence, potentially controlled (see definitions in Chapter 1, Section 7). In this prioritization model, object representations might be the *basis* of selection, rather than merely constraining spatial selection.

Much impetus for the prioritization model comes from work by Shomstein and Yantis (2002) who reported eliminating object-based effects on attention when target position was known with certainty. On the prioritization model, given known target position, there was no benefit for the subject to deploy attention to other parts of the object, so there was no observed object-based advantages to those parts. Chen and Cave (2008) have offered a different interpretation that is more consistent with the spreading account, but these matters remain an open empirical issue (for recent reviews, see Chen 2012 and S. Shomstein 2012; I'm grateful to Sarah Shomstein for discussion). Wannig et al. (2011) provided evidence that, as early as primary visual areas (V1), there seems to be an automatic spread of attentional modulation, namely, gain modulation, to neural representations of non-task relevant stimuli. Their results suggest that not only is there increased response to lines that are targets of attention, as recorded in neurons whose receptive fields are stimulated by that target, but also increased response to lines that are not targets of attention, but which stand in certain gestalt unity relations to the targets, namely, colinearity. No such effects were observed for task irrelevant lines when those lines did not stand in gestalt relations to lines that were targets of attention. Attentional modulation might automatically spread to task irrelevant objects if those objects are grouped to target objects, i.e., when the targets in question formed a larger gestalt group.

5 There are other demonstrative thoughts that are based on memory or testimony that we will not discuss. Note that we should distinguish between demonstrative thoughts and demonstrative expressions. On a classic and influential discussion of linguistic demonstratives, see (Kaplan 1985). For a pioneering discussion of demonstratives in thought that discusses this distinction, as well as non-perception-based demonstrative thoughts, see (Evans 1982).

- 6 There are many complicated issues here about the relation between thoughts and the sentences that we use to express them. I am assuming a type of isomorphism between sentences and thoughts that, for current purposes, is harmless. But it is easy to see that the same sentence type can be used to express different thoughts. Think of the sentence, "How I love you!" said sincerely and said ironically.
- 7 This is controversial, for one might think that a sortal is implicitly needed to lock on to the building, as opposed to its shape. The sortal helps to direct attention in an appropriate way. Still, the description also contains a proper name, and it is not clear that in understanding Jane's thought, one needs to invoke Jane. It may be enough to follow her gaze or her pointing. For now, I will fudge the issue by imagining that the relevant target drives attention bottom-up and independent of conceptual mediation.
- 8 This distinction resonates with Bertrand Russell's (1910) distinction between knowledge by acquaintance and knowledge by description.
- 9 Pylyshyn speaks of FINSTs as "nonconceptual" but it is important to note that when philosophers speak of the nonconceptual, they have in mind a certain kind of representational content, namely, one that does not involve concepts. When Pylyshyn speaks of FINSTs as nonconceptual, he is not thinking of them as representational. Rather, they are nonconceptual because they are not representational.
- 10 Which properties of tracked objects are used by a FINST to enable MOT? If the targets were visibly distinct, say, where targets are circles and distractors are squares, then distinct features could explain selective tracking (here, being a circle). But in typical MOT experiments, the targets and distractors are the same shape, and need not vary on any other visible feature. A second possibility is a target's spatial location, yet while location can explain how we initially individuate targets from distractors at the beginning of a MOT experiment, they do not explain individuation once objects begin to move, since targets will now continuously change locations. Another possibility is target trajectory. In one experiment, however, where objects disappear behind an occluder during MOT, tracking was observed to be better when the object reappeared at the location where it disappeared rather than at the location predicted by its projected trajectory (Keane and Pylyshyn 2006). This does not mean that trajectory isn't the basis of tracking (disappearance might trigger other mechanisms like memory for
last known location), but it does suggest that which properties causally sustain MOT is a complex issue. For Pylyshyn, the crucial point is that the constant element to which FINSTs are locked is an object and not its properties.

- 11 Sean Kelly (2004) and Mohan Matthen (2006) raise similar questions in discussing John Campbell's work.
- 12 See also (Evans 1982) for an important treatment of demonstrative thought where attention plausibly figures in the notion of an information link that allows for tracking an object. Campbell's work advances the debate in that it brings out attention explicitly, drawing on relevant empirical work.
- 13 While there are certainly empirical challenges to the causal efficacy of conscious experience, Campbell's response draws in part on an influential picture of the role of conscious vision in action proposed by Melvyn Goodale and David Milner (2004): conscious vision sets the targets for (unconscious) visual programming of motor behavior. Milner and Goodale's basic picture also involves attention: attentional effects on visual areas tied to consciousness yield conscious experience of those objects, where this experience can identify the target for unconscious visual guidance of movement. There are metaphysical worries about the causal efficacy of consciousness and, indeed, of the mental, even if Milner and Goodale are correct about the circuitry in the visual system. We set these metaphysical challenges aside. An important question is how the personal and subpersonal levels might appropriately interact, how conscious content affects mere information processing. In answer to this, Campbell speaks of a *commensurability* between personal and subpersonal systems. The link between the systems is achieved by what Campbell calls "topdown commensurability" or, to use ideas discussed in Chapter 2, goaldirected commensurability. In particular, conscious attention sets the goal for information processing by identifying a target. As the target is located in space, location can be used to bridge personal and subpersonal levels. This is not to say that spatial location is the entire story, but Campbell suggests that it provides a plausible way in which commensurability between the two levels can be achieved.
- 14 For Huang and Pashler (2007) "visual attention, in its most fundamental sense, is a selective visual process that governs access to consciousness" (599). This is a version of the gatekeeping view. The key concept in Huang and Pashler's theory is a *Boolean map*, which they argue is the "mechanism of visual access" and "corresponds to the information that can be consciously apprehended at one instant based on vision ... an observer's

visual awareness corresponds to one and only one Boolean map at any given instant" (602). A Boolean map is a spatial representation that partitions visual space between an area that is selected and one that is not. The Boolean map then corresponds to the selected area. The represented locations precisely cover relevant stimuli such as visible objects ("one should imagine the map as being ... shrink-wrapped to conform tightly to the object" fn. 2, p. 601). In other words, spatial representations are tied tightly to objects. Features are represented in Boolean maps, but a map can only represent one feature-value per dimension, e.g., one color, shape, or orientation. When a Boolean map flags a feature in this way, it is *labeled*. Huang and Pashler argue that (a) location is obligatorily encoded; (b) that only a single feature value can be selected at a time; (c) only a single feature value can be consciously accessed at a time, and yet (d) multiple locations can be accessed at a time.

- 15 Campbell's discussion emphasizes that conscious attention is a personallevel phenomenon, and yet the model he uses to explicate the role of attention in causing and guiding personal-level processes—namely, FIT plausibly concerns subpersonal processes. One might put the worry as a question: Does a *subject* really use feature maps, or is it rather a subpersonal system that uses these maps? If the latter, then one worries that the story does not explain how conscious attention justifies the subject's way of verifying propositions via the use of feature maps because it isn't the subject who uses feature maps. To draw on Campbell's relational view of experience, the subject uses the object and its features to verify propositions.
- 16 I am grateful to Dickie for making available to me drafts of her monograph on reference and demonstrative thought that is under contract with Oxford University Press.
- 17 Dickie's current work draws on additional empirical sources to explain the reliability of attention in delivering information about objects.
- 18 Dickie argues that the mind has a basic need to refer to particulars. It is an interesting question whether we have such a general need. A detailed argument for this claim is given in her manuscript of her monograph on reference and thought.
- 19 There is an analogous contrast in a disagreement between Ned Block and Jesse Prinz on cognitive accessibility. Prinz holds that attention is necessary to make perceptual representations accessible to working memory; Block, perhaps, can hold that perceptual representations of the right sort are already accessible to working memory, but that attention is needed for accessing those representations.

242 ATTENTION AND DEMONSTRATIVE THOUGHT

20 In the discussion of conscious attention and phenomenal salience, I noted something that is relevant here, namely, that the salience of an object might be explained by demonstrative thought about that object. In other words, it is the demonstrative thought that explains, in special cases, our sense of a distinctive phenomenal highlighting in attention. If so, such highlighting cannot also be the basis of demonstrative thought.

8

THE EPISTEMIC ROLE OF ATTENTION

8.1 Introduction

The final chapter focuses on the epistemology of attention, an area where there is need for much further discussion and much exciting work to be done. Two large topics are on the table: attention in justification and the role of attention in introspection. As philosophical work in these areas is in an early stage, the discussion will be more cursory than in previous chapters. The goal is to raise questions that point to areas where more exploration is needed. There are rich opportunities in the epistemology of attention.

Section 8.2 recalls the distinction between propositional and doxastic justification, and briefly considers whether attention is necessary and/or sufficient for the former. Section 8.3 then presents a general argument that attention is necessary for doxastic justification because it is necessary for the "epistemic action" of figuring out what to believe. This recalls the role of attention in agency, as discussed in Chapter 3. Section 8.4 turns to a possible counterexample to the necessity of attention in introspection, and makes some preliminary remarks about introspective attention of perceptual consciousness, that is attention used to inform introspective judgments about the phenomenal character of our perceptual experience. Section 8.6 then provides the Direct Model of Introspective Attention where the target of

attention is the phenomenal character of experience (thus, a form of introspective feature attention). Section 8.7 introduces the Transparency Model of Introspective Attention where attention that serves introspection is paradoxically directed outward to the external world. The advantage of this model is that it can appeal to a well-studied form of attention, namely, attention in perception, but there are significant challenges for the model, as discussed in Section 8.8. Section 8.9 then turns to elaborations of the Direct Model by Brie Gertler and David Chalmers. On their models, introspective attention is inwardly directed, and has some interesting properties. The question, then, is what the grounds are for positing the existence of this inwardly directed form of attention as distinct from the forms discussed in earlier chapters.¹

8.2 Attention and justification

Here is a basic case: normally, when I have a visual experience of a blue car passing by in front of my eyes, this leads to a belief that there is a blue car passing by me. Given that the circumstances are normal (eliminating hallucination, evil demons and aliens, tricks, strange beliefs, and the like), the resulting belief is based in an epistemically appropriate way on my experience. In part, this is because the experience provides evidence or reasons for my belief and in that way justifies it, and I form my belief on that basis. On this telling, perception plays a rational role in belief formation, namely, in justifying my belief.

Attention is important as well, for it seems that whether I form the belief about the car depends on whether I notice the car. For assume that I fail to notice it because attention is diverted elsewhere (though see Section 8.4). It seems likely that I will not then form the relevant belief about its color or indeed any belief about it—based on my perceptual experience of it. If so, my experience provides evidence regarding the color of the car, but given that my attention is directed elsewhere, I do not notice the evidence and do not take advantage of it. The same points can be made using "reason" instead of "evidence": my experience gives me reason for a belief which I do not respond to given my inattention (I will be somewhat lax about usage of "reason" and "evidence," using them interchangeably). So, it seems that attention might be needed to use evidence or available reasons to inform my beliefs. It is a further question, however, whether attention is needed for the experience to provide the evidence that it does. After all, a natural thought is that the evidence is there whether I notice it or not. Having evidence or providing reasons is one thing; noticing and using it is another.

The previous chapter invoked a distinction that is relevant in this context: propositional versus doxastic justification. Propositional justification can be characterized as follows:

Propositional Justification: if experience E provides reasons, evidence or support for p, then E propositionally justifies p.

(from Nico Silins and Susanna Siegel, forthcoming)

It is a small step to then speak of E as propositionally justifying a belief with the content p. As noted earlier, a subject might have an experience E that justifies the belief that p, yet not form that belief. So, the subject can have an experience that provides propositional justification without ever taking advantage of it. Doxastic justification occurs when the subject forms that belief appropriately on the basis of E: The belief that p is doxastically justified by E.² Doxastic justification of the belief that p by E implies propositional justification in place, one can then ask about the epistemic role of attention by considering its possible connection to propositional and doxastic justification.

Given the earlier point about noticing, it seems that attention plausibly has a role in doxastic justification, a matter to be discussed in more detail. Let us consider propositional justification and attention so as to set it aside. This is a topic that deserves greater consideration, and the only goal of the current discussion is to note cases that have been deployed to elicit intuitions about the epistemic role of attention. This allows us to get a general sense of the lay of the land. Here are some cases to pump intuitions. One might think that attention has little to do with propositional justification, but consider an intuitive case (and I emphasize, this is just to pump intuitions). Consider blindsight subjects who are prompted to report on a stimulus in their blindfield, say a vertically oriented line. Even if these subjects hazard correct guesses ("the line is vertical"), one might be inclined to deny that their blindsight states provide propositional justification for their guesses. If blindsight is phenomenally like seeing in the dark or literal blindness, then one might wonder how there can be vision-based propositional justification in such cases. After all, there is nothing visually there for the subject to respond to, no subjective sense of the world. Furthermore, even if one adds attention into the mix, as seems plausible, since

blindsight subjects are responding to a specific stimulus to inform their response, propositional justification still looks to be missing (recall the work with the blindsight patient GY discussed in previous chapters). If propositional justification was missing before attention in blindsight, it is not secured simply by adding attention. Assuming that this gloss on blindsight conforms with your intuitions, it suggests that attention would not be sufficient for propositional justification in blindsight. Blindsighters select information from their blindsight states to yield accurate guesses, but those states do not provide propositional justification. Again, one might have different intuitions, but let us develop this line of thought, for, as seen in the last chapter, blindsight has been used as a device to test intuitions about attention.

What if one adds consciousness into the mix but without attention? This recalls John Campbell's sea of faces case discussed in Chapter 7: a subject is surveying a sea of faces trying to locate the person being talked about. Although the subject cannot lock attention onto the person, at least in a phenomenal sense of locking on, oddly, the subject can provide accurate reports about the person. In one sense, the subject is like the blindsighter, for the accurate reports are guesses in that the subject cannot articulate the reasons or grounds for her reports about the targeted person (although recall that the subject is not guessing in the sense of giving a report on the basis of no information, as if her eyes were closed. See Chapter 7, Section 7.6.3). In another sense, the subject is different from the blindsighter, for the subject is visually conscious, with no region of the visual field in darkness or phenomenally blinded. Does the subject in the sea of faces case have propositional justification for the resulting accurate reports, even though the reports are guesses? Here, intuitions might not be as clear. Some might hold that the subject has propositional justification since perception is conscious. This is the ingredient missing in blindsight, and once it is restored, perception-based propositional justification is available. The problem for the subject in the sea of faces case is that she cannot use perception in doxastic justification, for attention is defective. Others might hold that propositional justification is not present because attention is not functioning appropriately. If so, then attention might be necessary for propositional justification.

For myself, I find it hard to have strong feelings either way as to whether attention is necessary or not for propositional justification. Perhaps the same will be true of you. Still, it might be hard to accept that attention is necessary for propositional justification. Perception is one thing, attention another, and it is perception that provides propositional justification. If propositional justification is a property of perceptual experience, then it does not necessarily depend on attention, but is fixed by the subject's perceptual state. Nevertheless, there are arguments that suggest two ways in which attention is necessary for propositional justification: (a) attention might be necessary for there to be propositional justification at all; and (b) attention might be necessary for specific instances of propositional justification, say, an experience propositionally justifying a specific belief that p. Thus, one argument concerns a general requirement on propositional justification; the other a requirement on specific instances of propositional justification.

An argument for the general case relies on the gatekeeper view. If conscious perception is necessary for propositional justification, then if the gatekeeper view is true, attention would be necessary for propositional justification, since attention would be necessary for conscious perception. An experience can have the property of providing perceptual justification as such only if the subject is attending to the world so as to yield perceptual consciousness. Thus, gatekeeper theorists might endorse the view that attention is necessary for perception to provide reasons for belief in propositional justification.

The argument that ties attention to specific propositional justification of a belief is different, and it is less clear that there are philosophers who endorse the resulting view. To construct the argument, one must cobble together different theoretical claims. Some philosophers hold that propositional justification, i.e., perception's providing reasons for belief, depends on perception being conceptual in the sense of having conceptual content (McDowell 1994; Brewer 1999). The background assumptions are that beliefs have conceptual contents if any mental states do (concepts are the building blocks of thoughts), and only conceptual contents can serve as reasons for belief states. On this view, if perceptual states can serve as reasons for beliefs, then perceptual states must have the same kind of contents as beliefs, namely conceptual contents (this is a controversial claim; for an early response, see Peacocke 1992). This seems to intellectualize perception in the sense of holding that the contents of belief and experience are of the same kind, so creatures without concepts, and thus the capacity to have conceptual thoughts, cannot have perceptual experience of this concept-involving kind.3

Let us assume then that perception has conceptual content. If one also holds that the relevant perceptual conceptual content, say, the content of a visual experience that that car is blue, is available only if one attends to the car and its color, perhaps because attention must be focused to bind the relevant concepts together in perception, then it would follow that attention is required for specific cases of propositional justification (this adapts some ideas on attention and conceptual binding gestured at in Smithies 2011b).⁴ Thus, take the visual experience that the car is blue, an experience that provides reasons for believing that the car is blue. On the current view, attention to the car and perhaps its color is necessary for perception to have the conceptual content that the car is blue. Only then is there propositional justification of the relevant kind for the belief in question, namely, a belief that p that is propositionally justified by the experience that p. It is not clear, however, that anyone holds this specific view even if its various components are held by different philosophers. Still, these are certainly matters requiring further exploration. Epistemologists are at the beginning of sustained philosophical discussion on these matters from the perspective of a serious engagement with attention.

8.3 Is attention necessary for doxastic justification?

The answer here seems obviously, "Yes," given the earlier example of noticing: to form a belief on the basis of experience, one must notice the relevant aspects of experience. This suggests a more general way to argue for the necessity of attention for doxastic justification, one that draws on our discussion of attention and agency in Chapter 4.⁵ Let us begin with an observation by John Turri:

In evaluating beliefs we are evaluating a kind of performance, the performance of a cognitive agent in representing the world as being a certain way, and when performing with materials (which, in cognitive affairs, will include reasons or evidence), the success, or lack thereof, of one's performance will depend crucially on the way in which one makes use of those materials.

(Turri, 2010, p.315)

Certainly, in many cases, this description seems exactly right. If belief aims at truth, then fixing belief involves figuring out the truth. Figuring out the truth is often a cognitive activity, something done intentionally, in a goal-directed way. How effectively this activity is carried out depends on the manner of execution and on the materials at hand. To focus on perception-based beliefs, good epistemic standing depends on whether experiences provide propositional justification and whether beliefs are formed in a way such that they are doxastically justified. In aiming for the truth, one often has to do two things: searching for relevant evidence or reasons, say, by looking or otherwise experiencing the world, and considering and weighing the evidential options. Will it rain in the afternoon? One could, of course, hazard a guess, but to know the truth, one will look for and assemble relevant evidence and, on the basis of that evidence, form a belief about the weather. For example, one can ask a knowledgeable friend or just look outside.

For mundane cases of figuring out the truth, there is a cognitive action, and where there is action, the Many-Many Problem will apply. One simply has to delineate the relevant behavior space, in this case an epistemic behavior space. Specifically, the space is structured by inputs as sources of evidence for belief and the outputs as beliefs. Generally, there are many kinds of sources of evidence or reasons: experience, beliefs, memories, and such. I will focus on a subpart of a typical epistemic behavior space, namely, where the inputs are perceptual experiences and the outputs are beliefs that could be caused by those experiences. There are then three experience-belief mappings to consider. The simplest mappings link experiences and beliefs that have contents with the same accuracy conditions, e.g., the belief and the experience that p. More complicated mappings link beliefs whose contents are propositionally justified by the available experiences, a category that will include more beliefs than in first case where experience and belief have the same content. Finally, there will be beliefs that are only apparently propositionally justified by the set of available experiences. That is, a subject might falsely think that certain beliefs are propositionally justified by a certain experience and form those beliefs on that mistaken basis. Adding this third type of belief into the mix yields a complicated behavior space where experiences are connected to: (a) beliefs with which they share the same content; (b) beliefs for which the experiences provide actual propositional justification (which include (a)); and (c) beliefs for which they provide apparent propositional justification, from the subject's perspective. Of course, the space is more complicated than this, but the current layout suffices for our discussion. The epistemic Many-Many Problem then defines a space of possible perception-based beliefs. That there are so many possible paths is driven by the idea that the path one takes depends on what one notices (i.e., attends to).

"Base" here is used in a broader sense than it is in epistemology where much focus has been directed at the basing relation. Ram Neta writes: [The basing relation] is that relation *the obtaining of which* makes it the case that the reasons that stand on one side of the relation are *the reasons for which* the creature holds the belief that stands on the other side of the relation. In other words, the basing relation is that relation which is such that, when it obtains between a belief *B* and a reason *R* for which the belief is held, its obtaining is what makes it the case that *R* is the reason (or at least *a* reason) for which *B* is held.

(Neta, 2010, p.109)

If one thinks that visually experiencing that the car is blue provides evidence for one's belief that the car is blue, then, where believing that the car is blue is based in the relevant way on experience, this is a case where the perceptual reason is one's reason for that belief. In this case, the belief is well-founded. Here, the evidence propositionally justifies belief and the relevant basing relation holds, so the belief is doxastically justified. The more general notion of basing as used to define the epistemic behavior space includes the epistemic basing relation as among the mappings linking perception to belief. It is one among many ways to form one's beliefs. In the epistemic behavior space, there are three salient cases for experience E and belief B where B is based (in the general sense) on E:

- 1. E propositionally and doxastically justifies B (i.e., B is properly based on E);
- 2. E propositionally, but does not doxastically, justify B;
- 3. E does not propositionally justify B, yet B is based on E.

Recall that the space characterizes the goal-directed activity of figuring out what to believe. Since beliefs aim at the truth, one figures out what is true in figuring out what to believe. It is natural in this context to "look" for specific evidence or good reason, and in this way, interrogate one's putative evidential sources selectively. When one finds relevant evidence, one focuses on it, considers it, and evaluates it in respect of other evidence. Attention is part of this activity of focusing, interrogating, and considering. Call the epistemic behavior space at issue the space of experience-based belief. The argument for the necessity of attention for doxastic justification is just the argument for the necessity of attention for beliefs based on experience in the cognitive activity of figuring out what is true. As doxastic justification is one of the three forms of basing noted previously, the necessity of attention for doxastic justification for doxastic of attention for cognitive action. To figure out what is true is an activity that culminates in traversing a specific path in behavior space, and the selectivity at issue, namely, the specific belief formed in light of a specific perceptual experience, implicates attention. Think of attention here as selection of perceptual content to inform the fixation of belief. So, in the cases of cognitive activity for the sake of determining what to believe, attention is necessary for basing belief on experience in the general sense of basing, and *a* fortiori attention is necessary for doxastic justification where a specific epistemic basing relation is required, whatever that relation turns out to be. It follows then that the epistemic basing relation will involve attention.

In light of this, one can now bring the distinctions discussed in earlier chapters on board to formulate specific questions:

- Does doxastic justification entail a specific form of attention, say topdown/bottom up and/or controlled/automatic attention in respect of the relevant experience?
- Does doxastic justification entail specific targets or combination of targets of attention, say, of an object, feature or space in respect of the relevant experience?
- Does doxastic justification entail conscious attention in respect of the relevant experience?

These are still fairly general questions, but the point is that, to the extent that attentional selectivity has a role to play in doxastic justification, there is much work to be done to sort out the precise role that attention plays. These are open questions that need answers. To delve into this a bit further, let us now consider a counterexample to the necessity claim.

8.4 Might attention not be necessary for doxastic justification?

Nico Silins and Susanna Siegel (forthcoming) have raised doubts about whether attention is necessary for doxastic justification even if it often serves this role. Here is one of their cases:

Consider a distracted subject navigating the environment, such as a distracted driver, or a walker lost in thought. Such a subject can still adjust their behavior in response to the environment in a way that is not merely instinctive, operating the brake, the clutch, the defrosting system, the steering wheel, and so on. Further, they arguably can do so while remaining distracted from the environment, without their attention being captured by the obstacles that they are successfully avoiding, or by the equipment that they are manipulating. Despite being superficially automatic, such behavior is far from being a mere reflex, and has a strong claim to being rational. In addition, the subject would satisfy the central diagnostics for having various beliefs about her immediate situation, such as the belief that the car is running and operating as it should be. They are disposed to endorse this proposition if asked, and they are acting in a way that would be advisable, given their desire to continue driving, if the proposition is true. In such a case, their inattentive experience is feeding into well-founded perceptual beliefs.

In presenting this example, Silins and Siegel are targeting a specific phenomenal conception of attention, similar to experiential highlighting (see Chapter 4). Since the previous section focused only on attentional selection as a functionally characterized notion, we can now raise an issue that was broached in earlier chapters: Is the required form of attention a functional kind or a phenomenal kind? One might hold that conscious attention is required in order to fix belief. Silins and Siegel's example suggests that attention in the phenomenal sense is not necessary for doxastic justification. One can have well-formed beliefs about items that are not phenomenally highlighted in experience. If this is correct, then doxastic justification does not entail conscious attention. This is in line with concerns raised in the last chapter regarding the necessity of conscious attention for demonstrative thought.

Let us push their argument further and ask whether it shows that attention is not necessary in doxastic justification.⁶ Much rides on what inattention ultimately comes to, and this ties in with our earlier, brief discussion of partial attention (Chapter 3). In characterizing attention as selection for action, attention has a necessary connection to coupling. To the extent that there are multiple instances of coupling in performing a task, i.e., multitasking, one can quantify degrees of attention in terms of degrees of coupling. Thus, in multitasking, a subject might select multiple items for action, and in inattention, select fewer items. Alternatively, one can capture inattention by drawing on the contrast between control and automaticity, with attention correlated with the amount of control, inattention with the amount of automaticity. In this case, distraction is where attentional control decreases. One can also discuss inattention in terms of the distinction between conscious and unconscious processes, and, like the control/automaticity case, ascribe increasing inattentiveness to decreasing consciousness. There may be other possibilities, but this will serve as a start.

The question is whether doxastically justified beliefs about an object require attention to that object.⁷ The distracted subject in the previous counterexample, however, need not fail to attend to relevant objects, even if her experience of that object is inattentive. She can still be attending to the object even if she is thinking about something else. Given that thought about X can entail (cognitively) attending to X, the fact that the subject is lost in thought about something else is just a way of being less attentive to the environment. If one characterizes inattention in terms of coupling, then there is one less selective process deployed in perceptual selection and one more selective process directed elsewhere; on the control/automaticity view of inattention, it is less control in perceptual selection, with control shifted towards thought; and in the conscious/unconscious view of inattention, it is less conscious selection of the environment, with conscious selection apportioned to thought. If one of these approaches is adequate, then there is an explanation of attending to an object inattentively. Yet inattentive attention is still attention.

The current line of thought suggests that while Silins and Siegel's counterexample might show that conscious attention is not necessary for doxastic justification, this leaves untouched the idea that attentional selectivity with respect to perceptual evidence is necessary for doxastic justification. This touches on one of the three questions posed at the end of Section 8.3: Does doxastic justification entail conscious attention in respect of the relevant experience? The present answer is, "No". The main point, however, is that once attention's role is clear, a set of interesting questions about doxastic justification emerge. There is much work here to do in clarifying the epistemic role of attention in justification.

8.5 Introspection and attention

In the remainder of the chapter, I shall consider attention in introspection, a distinctive source of knowledge about one's own mental states. Introspectionbased beliefs seem in many ways epistemically privileged. One traditional way of understanding privilege is that introspection-based beliefs are certain or infallible. Famously, Descartes made much of this feature in his *Meditations*, but these days, philosophers are much less inclined to take introspection-based beliefs as infallible except in very specific cases. Rather, philosophers agree that introspection-based beliefs are fallible, yet are typically epistemically superior to perceptually-based beliefs (but see Schwitzgebel 2011 for a dissenting view).⁸ My introspection of my pain state that leads to my belief that I am in pain seems to put me in an epistemically privileged position vis-à-vis my pain, as opposed to your belief that I am in pain formed on the basis of your observing my behavior. There are many complicated issues regarding the epistemic privilege that introspection possesses over perception (see Gertler 2011 for a comprehensive overview of self-knowledge). Our focus will be on attention as it figures in introspection, and specifically in introspection of ongoing perceptual experience.

The challenge is to explain how attention works to secure a demonstrative belief about a specific phenomenal feature of ongoing perceptual experiences, say, the phenomenal feature associated with the experience of color, of pain, or of the pitch of a sound. This is an underexplored topic, and yet attention seems central to introspection-based beliefs about consciousness. Let us focus on two models of attention in introspection, or as I will refer to it, introspective attention: (1) introspection that involves attending to the inner, mental world; and (2) introspection that is (paradoxically) achieved by attending to the outer, external world.

Plausibly, introspection depends on a form of attention that enables selective thought about mental properties. Introspective attention is attention that grounds introspective thought, that is thought about the mental. This capacity is a familiar one. Consider Frank Jackson's (1982) story about Mary who, on being released from a chromatically colorless room, sees red for the first time and knows immediately what it is like to experience red. That she can focus on the specific phenomenal feature associated with experiencing redness, as opposed to familiar phenomenal features associated with experiencing shape, texture, or achromatic colors, implicates a selective grasp of that feature. This suggests that she attends to a specific phenomenal property among the many that her visual experience exemplifies. In general, to think introspectively about a phenomenal feature, one has to select that feature against many other such features. Call this the Many Phenomenal Features Problem, a version of the Many-Many Problem (Chapter 3). Introspection relies on attentional selection that yields a solution to this problem.

As with any form of attention, introspective attention can be characterized in terms of its inputs and/or its outputs. Let us start on the input end with the target of introspective attention. If introspection relies on attention, then the relevant form of attention is naturally individuated from other forms by its distinctive target, namely, the mental. More specifically, it might be a species of *object* attention, as when one introspects mental events like an experience, or a species of *feature* attention, as when one introspects the phenomenal character of those experiences or their intentional properties. As I shall focus on introspection of phenomenal character, I shall focus on introspective attention as feature attention. But consider now the output end. At least part of what makes this form of attention *introspective* is that it is geared towards enabling thought about the mental. The two models of introspective attention to be considered agree about the output of introspective attention in that attention informs introspective thought; they simply take attention to be directed in different ways on the input side.

The distinctions applied earlier to attention are applicable here. For example, introspective attention can be automatic or controlled, bottom-up or top-down (see Chapter 1). Introspective attention can be controlled and top-down, as when one attempts to ascertain facts about one's mind and deploys attention as a way to answer questions about it: Is current vision blurry, or is it that the object is merely fuzzy? Here, there is an epistemic action, figuring out what is right in respect of the state of one's mind. Attention can also be automatic and bottom-up, as when a pain drives one to certain thoughts about it ("gosh, not that again") or when Mary cottons on to a previously unknown phenomenology of chromatic color experience, say the phenomenal character associated with experiencing red ("that's what it's like!").

It is important to emphasize the possibility of bottom-up and automatic attention. Sometimes, it is claimed that introspective attention must be top-down in the sense of being cognitive. Thus, David Chalmers writes:

As with all acts of demonstration and attention, phenomenal demonstration and attention involves a cognitive element. Reference to a phenomenal quality is determined in part by cognitive elements of a demonstration. These cognitive elements will also enter into determining the content of a corresponding direct phenomenal concept.

(Chalmers 2003, p.237)

Ernest Sosa speaks of selective attention in introspection as "the index finger of the mind" and then comments that:

Of course the requirement of such a mechanism of attention would seem to import a need for some presupposed non-demonstrative content to the attention, which means that such demonstrative reference could not be conceptually fundamental.

(Sosa 2003, 279, fn. 5)

These claims might be perfectly correct for top-down introspective attention, but there might nevertheless be a bottom-up form of introspective attentional capture that yields a conceptually simple demonstrative thought about that (a phenomenal feature). Think of a sudden strange, alien feeling that one does not know how to categorize, a feeling that suddenly enters consciousness and drives the question: What is that?

Let us now consider two ways of unpacking the notion of introspective attention. It is fairly uncontroversial that introspection involves attention to phenomenal features at least in the sense that one can deploy a selective capacity to fix a thought about those features. But what exactly is this selective capacity?

8.6 The Direct Model of introspective attention

When one introspects a current visual experience, one can, on the basis of introspection, form a judgment or otherwise think about the phenomenal properties of the experience. A natural model of introspective attention is the Direct Model:



Figure 8.1 The Direct Model of Introspection.

Introspective attention, represented by the arrow, is what allows the subject to form an introspective judgment about a specific aspect of experience. The arrow can be understood to indicate a causal process, or one can leave it open that there might be non-causal ways of unpacking the contribution of introspective attention (see Section 8.9). Furthermore, one can think of attention as occurring at two stages, first in introspective attention that grounds the introspective judgment, or in the introspective judgment itself, construed as a form of cognitive attention, a thought about a specific phenomenal property. This judgment then constitutes the subject's awareness of the phenomenal property.

The previous model is "basic" in the sense that it captures the idea of an internally directed capacity that grounds internally directed thought. Here, attention's input and output are focused inwards: one attends to a specific phenomenal character of experience and as a result one thinks about it.

8.7 The Transparency Model of introspective attention

There is another possible model of introspective attention. Let us begin with an observation by Gilbert Harman:

When Eloise sees a tree before her, the colors she experiences are all experienced as features of the tree and its surroundings. None of them are experienced as intrinsic features of her experience. Nor does she experience any features of anything as intrinsic features of her experiences. And that is true of you too. ... Look at a tree and try to turn your attention to intrinsic features of your visual experience. I predict that you will find that the only features there to turn your attention to will be features of the tree.

(1990, 667)

Harman voices what I will call the transparency observation: attempts to introspect phenomenal properties in the sense of intrinsic features of experience lead attention to the external world. The transparency observation has been used as the basis of an argument against qualia theories of phenomenal consciousness or in favor of representationalist theories of consciousness (see especially Tye 2000). There are genuine questions, however, whether the transparency observation can support such arguments (for critical discussion see, among others, Kind 2003, Martin 2002, Molyneux 2009, Siewert 2004, and Stoljar 2004). I shall, however, sidestep these broader debates and focus on something explicitly raised in Harman's transparency observation: the role of attention.

Note that Harman makes a prediction: "Look at a tree and try to turn your attention to intrinsic features of your visual experience. I predict that you will find that the only features there to turn your attention to will be features of the tree." The claim is that attention locks on to the external world. A sense-datum theorist, however, will remain unimpressed, for they will deny that one is thereby directly aware of the external world. For all one can tell, the sense datum theorist continues, attention is locking on to sense data or internal mental objects that one mistakes for something external. I think this is a plausible initial retort, but rather than getting bogged down on this issue, I assume realism in the sense that attention does lock on to the world (one can treat this as a hypothesis).

The transparency observation then suggests the Transparency Model of introspective attention of perceptual experience



Figure 8.2 The Transparency Model of Introspection.

Recall that "introspective attention" means attention as used in introspection. Given the Direct Model of introspective attention, which I take to be the intuitive one, the Transparency Model is radical. On the Transparency Model, the target of introspective attention is not perceptual experience, but the external world. Thus, introspection is in a sense outwardly directed in that it depends on perceptual attention. Yet introspection is, definitionally, inwardly directed. If introspection involves introspective attention, then an inwardly directed capacity relies on an outwardly directed one, perceptual attention. This seems paradoxical. The fundamental difference, then, between the Direct Model and the Transparency Model is that while they agree that attention informs introspective judgments, they disagree about the inputs to attention, i.e., the targets of attention. The Direct Model takes the target of attention to be the mental; the Transparency Model takes it to be the external world.

A pressing question is how attention to the external world can support judgments about experience. Still, a route from the external world back to the mind is readily available. One's experiences are intentional and thus about the world. In attending to the world through experience one gains information regarding what the experience is about. That is, one gains information about the intentionality of experience. If the Transparency Model is correct, then one accesses the phenomenal character of experience by accessing what one's experiences are about. Where perceptual content identifies what is experienced, attention provides reflection with the content of experience and, in that way, points to an aspect of what experience is like.

The challenge then is to explain how perceptual information that is normally selected to inform thought about the external world gets detoured to inform thought about the internal world. But this need not be difficult to understand, at least in outline. In forming a thought on the basis of perception, one brings concepts into play, namely, whatever counts as the building blocks of thought contents. When one thinks about the external world, one deploys "empirical" concepts, i.e., concepts about the external world. These concepts then engage with the deliverances of perceptual experience to provide for perception-based thoughts. Similarly, introspection yields thought about the inner world, and if the deliverances of perception about the external world have some relevance to our understanding of the mental world, specifically in respect of the intentionality of experience, then again, perception need only engage appropriate concepts, in this case concepts of the mental. Thus, one will call on psychological concepts, including concepts regarding experience.

The major challenge to the Transparency Model, then, is not how perception of the external world provides material of relevance to our understanding of the mental world, but rather to explain how one acquires the specific concepts of the mental that one uses to think about consciousness, since acquiring these concepts might also require introspective attention (see Section 8.9). Since these concepts are about the mental, it would then seem that introspective attention, at least in some of its forms, has to be internally directed. Thus, even if the Transparency Model is true for many cases, fundamentally, the Direct Model is basic. The latter is needed to explain the acquisition of psychological concepts of experience that the former simply assumes.

The Transparency Model seems congenial to representationalist accounts of phenomenal character (see Chapter 4). For, on those views, content is or determines phenomenal character, so if the Transparency Model shows how to attend to content, i.e., to what one experiences, it thereby explains how to attend to phenomenal character. There is a slight complication for those representationalists who do not identify content with phenomenal character since a second step might be needed to move from content to phenomenal character, but set that aside. Transparency also seems congenial to relationalists about experience, since on that account, phenomenal character is constituted by the perceived properties of objects, so attention to those external features is just a way of attending to what constitutes the phenomenal character of experience.

I began with the transparency observation which, taken at face value, supports a Transparency Model of introspection. There are reasons to worry whether the model is viable, but let us note one positive reason to accept it: it renders introspective attention relatively intelligible, for introspective attention just is perceptual attention used in a distinctive way to inform thought about the mental, and there are detailed philosophical and empirical accounts of perceptual attention (see Chapters 1–3). The Transparency Model is thereby a serious account of introspective attention despite its somewhat paradoxical stance that one attends to the inner by attending to the outer.

8.8 Challenges for the Transparency Model

Things are, of course, not that simple. If introspective attention just is perceptual attention deployed in a different way, then an account of the former depends on an account of the latter. But the identity claim raises a host of questions. First, introspective knowledge is regarded as typically epistemically superior to perceptual knowledge, yet the Transparency Model seems to obliterate that advantage since it identifies introspective attention with perceptual attention. At best, introspection looks like it can be no better than perception. This is a significant challenge that must be addressed, but I shall continue to focus on the nature of introspective attention by raising a different challenge.

Sometimes one has vivid hallucinations, experiential states that clearly have phenomenal character yet where there is no corresponding object in the world. When visually hallucinating pink elephants, there are in fact no pink elephants as one experiences there to be. Accordingly, there is no appropriate external object to attend to. Yet no one doubts that one can introspect hallucinatory experiences reliably and accurately, and in such cases, it seems that one can have direct knowledge of the nature of those experiences by introspection. Introspection must lock onto the pinkness of the elephant, yet there is nothing that is actually pink to attend to! What then is the target of introspective attention in hallucination such that it can lead to knowledge of the features of hallucination?

This is a version of a challenge to relationalist accounts of perceptual experience, namely that they cannot explain the phenomenology of hallucination (or illusion).⁹ After all, relationalists account for phenomenal

character by focusing on external objects and their properties, but these are absent when the subject hallucinates. Representationalists might feel more confident here, for their point is that qualitatively identical veridical experiences and hallucinations share the same intentional content. So, one can access the shared phenomenal character of hallucinations and veridical experiences by attending to content. But can one attend to content? Recall the simple notion of content as what one perceives. This notion fits seamlessly with the basic transparency observation: what one perceives is the world attended to. Yet for representationalists, content is a technical notion that refers to an abstract entity, something that determines an accuracy condition. Such entities are variously conceptualized as Fregean senses, Russellian propositions, or sets of possible worlds. We need not discuss the details behind these options, but simply note that whichever of these is chosen as the contents of mental states, the object of attention will thereby be an abstract entity. Yet the perceptible world is not abstract but concrete. This shift from a simple to a technical notion of content threatens to spoil the transparency observation. It is not immediately clear how the Transparency Model can explain introspection of hallucination.¹⁰

The advantage of the Transparency Model of introspective attention is that it can appeal to a rich body of work on perceptual attention that fleshes out its account of introspection. Yet that work treats attention as directed to the world. In light of hallucination, the Transparency Model of introspective attention threatens to become disjunctive, telling one story for veridical perception (attention to physical objects and their instantiated properties) and another for hallucination (attention to abstracta or uninstantiated properties, if to anything at all). Once attention to something not in the world is on the scene, one might just wonder why not bite the bullet and opt for a uniform account? It is content in a technical sense (contents that determine accuracy conditions) that is a common factor in the perceptual states one introspects, whether they are veridical experiences or hallucinations. So it is content that is attended to in introspecting both types of experience. Such a view essentially jettisons the transparency observation that ties attention to the physical world, and if it is jettisoned, why not opt for a more basic view: in introspection, one directly attends to a phenomenal property, an internal feature of experience. From the point of view of good sense, that can be no worse than attending to contents in the technical sense, and it has the advantage of providing for a uniform account: in both hallucination and veridical experience, introspective

attention is locked onto the same thing, phenomenal character. This returns us to the Direct Model of introspective attention.

8.9 Direct attention to phenomenal character

The challenge to proponents of the Direct Model is that they need a story about introspective attention, for they assert that one can attend directly to phenomenal character. Brie Gertler (2001) and David Chalmers (2009, chap.8–10) have presented accounts of introspective attention.¹¹ For both, attention allows one to think about those phenomenal features in a unique and immediate way.

Gertler spells out her account in terms of phenomenal state introspection (PSI):

S introspects her occurrent phenomenal state token a with phenomenal character F, iff S has an occurrent mental token b which is such that:

- 1. *a* is embedded in *b*;
- 2. b refers to F; and
- 3. 2 is true partly in virtue of 1.

(Gertler, 2001, p.307)¹²

This provides a metaphysical and semantical framework for introspection that can be spelled out in more explicitly psychological terms by adding attention: When a subject attends to an experience a exemplifying phenomenal feature F, so as to form a judgment b about that phenomenal character, then the judgment *embeds* the experience as a part. So, if I attend to my visual experience of a red rose and specifically to the visual phenomenal character associated with seeing red, the resulting judgments about that experience can embed the exemplification of the phenomenal character as a part. A complex mental state results where the judgment refers to the embedded phenomenal character. Gertler also emphasizes that it is embedding that enables direct reference to F (hence, clause (2) in PSI might be helpfully rewritten as "b directly refers to F," since the embedding also explains the referential directness).

What is the significance of "direct"? In part, directness in light of embedding emphasizes a distinctive first-personal aspect of introspective attention. For example, one might make judgments about another person's pain states by neuroimaging, looking at the person's brain activity. Similarly, you could image your own brain, and in that way get evidence regarding your mental states. Gertler considers a case where a brain scan reveals that one is in a neural state N which, for argument's sake, is necessarily correlated with mental state M. Accordingly, one can on the basis of the neuroimage form the judgment that one is in M. This judgment is clearly formed in an indirect way, one that depends on perception of the image. There's nothing special or uniquely first-personal about such access, as an external observer has the same kind of access to the mental state via neuroimages. If such access is indirect, one can define direct access negatively in terms of a contrast with the indirect route to M provided by perception.

It would, however, be better to provide a positive characterization of direct access. A natural response would be to invoke attention as oriented directly at the relevant internal property, but recall that the current challenge is to say more about the nature of attention here. Gertler's positive proposal is to appeal to embedding. When a judgment embeds its referent in light of attention, this allows for a distinct relation between the judgment and what it refers to such that reference failure is impossible. That is, if reference to some phenomenal feature is a result of embedding it, the phenomenal feature referred to cannot fail to be there. In that sense, reference failure is impossible. This provides a contrast to perceptionbased, indirect forms of reference, as in the neuroimaging case. Embedding allows for a contrasting direct form of reference.

Crucially, embedding is not a causal relation, but secures reference by constitution, i.e., by making the referent literally part of the referring state. In this way, embedding makes available direct ways of thinking about phenomenal features. Gertler conceives of the resulting phenomenal concepts as demonstrative concepts. In Chapter 7, I discussed perceptual demonstratives and the role of attention in securing demonstrative reference. Phenomenal demonstratives differ from perception-based demonstratives in that attention is directed at phenomenal features. Yet while both types of demonstratives depend on attention, there is, again, a critical contrast. Many perceptual demonstratives require a descriptive component to secure reference, say, an accompanying sortal term (that cow):

When we perceptually demonstrate something, we use an implicit descriptive component so long as we presume that the referent must be something to which we stand in an appropriate causal relation. And this presumption is always present in perceptual demonstration.

(Gertler 2001, 315)

Phenomenal demonstratives, however, secure reference independently from either descriptive mediation or the use of ordinary means of demonstration, such as pointing. Rather, attention to phenomenal features is sufficient for securing demonstrative reference to those features. Gertler speaks of these demonstratives as pure (phenomenal) demonstrative concepts.

Chalmers' account shares many of the same structural features as Gertler's.¹³ One possible difference is worth mentioning in this context. Chalmers takes attention to explain embedding while Gertler does not, at least not explicitly (Chalmers suggests this reading in his 2003, 169, fn 15). On Chalmers' account, embedding (clause (1) in PSI) and reference (clause (2) in PSI) have a common basis: introspective attention. So, PSI can be fleshed out in light of Chalmers' proposal of the role of attention in this way: attention to a phenomenal feature is sufficient for embedding it in a judgment, a constitutive relation that allows for a distinctive sort of direct reference to that feature, one where failure of reference is not possible. Chalmers refers to the resulting phenomenal concepts as direct phenomenal concepts.¹⁴ I think Chalmers' elaboration is a reasonable proposal of the function of introspective attention in securing reference for phenomenal concepts, and since it is compatible with Gertler's framework, I shall assume that both Gertler and Chalmers accept that introspective attention explains embedding.

Gertler and Chalmers fill out a version of the Direct Model that has many striking features, but I shall keep focus on the nature of introspective attention. I have been contrasting the Direct Model with the Transparency Model, and it seems that the capacity for introspective attention differs between the two. That is, each posits a distinct psychological capacity for attention. On the Transparency Model, attention is directed to the external world; it cannot literally be directed inward. Furthermore, attention in transparency is not sufficient for embedding. It does not, after all, embed external world objects or features into any resulting judgment.¹⁵ Challenges to the Transparency Model aside, it has the virtue of appealing to a fairly well understood psychological capacity: perceptual attention. On Gertler and Chalmers' version of the Direct Model, introspective attention can directly lock on to phenomenal features and, in doing so, embed those features in judgment. There is something intuitive about their version of the Direct Model, but since they posit a different type of attentional capacity, one with different functional properties, what more can be said about it? Gertler raises the possibility that the notion of attention here is a primitive one, not further analyzable (she does not necessarily endorse this option).

Perhaps that is true, but the question I want to raise is this: Can one even isolate this capacity?

Harman's original observation assumed that one can intentionally deploy introspective attention, the capacity currently at issue. That seems relatively uncontroversial: one has enough of a grasp of the capacity to use it voluntarily. But now consider a challenge inspired by Harman's observation. First, visually attend to some object in the environment. Clearly you know how to do that. Is the item you are attending to hot pink? If not, then think about its color and how it relates to other colors (or whatever thought tickles your fancy). Clearly, thinking about the color and perceptually attending to it invoke two different capacities. You can tell these capacities apart, and you know when you are deploying one rather than the other: focus on the attending and then on the thinking. Both amount to two different ways of attending to color, one perceptual, the other cognitive.

Now focus again on the color of the object you are attending to. Notice its shape, texture, and location. Now shift your attention to the corresponding phenomenal character for each external feature you were just visually attending to. Shift attention across the many phenomenal features your experience exemplifies. Something different is happening when you attend to the external world and then shift to the inner, but here is the challenge: What is the difference? There is a shift of attention, but what if the shift is just in the output of perceptual attention, namely, toggling between thought about the world to thought about the inner, and not in the nature of the selective process that informs thought. This shift in output is a shift in cognitive attention from being directed at the external world to being directed at the internal. That is, the shift from perceptual to introspective attention is just a shift in thought. This is a version of the transparency observation, here applied to attention rather than the objects of attention. The point is that this shift in the output of attention might suffice to capture the idea of introspective attention to the inner as advocated by the Direct Model. If the Direct Modeler disagrees, then the question for them is this: Where in all the various deployments of attention in the previous exercise is the form that suffices for embedding?

Let us be clear on the nature of the challenge to the Direct Model. It begins with the natural idea that one can attend to the mind. Yet confronted with a competitor that identifies introspective attention with perceptual attention geared towards an introspective thought, the Direct Model was revealed as inadequate: it had no story to tell about the nature of attention beyond the natural idea. Gertler and Chalmers do better, but the capacity of attention extracted from their discussion is a striking one, namely, a capacity that embeds the phenomenal feature in the resulting judgment. The previous challenge was simply that while one can recognize the deployment of cognitive and perceptual attention, one does not discern a third form of attention, that capable of embedding. I think this is puzzling.

It is not clear that we have as yet good reason to postulate the existence of introspective attention of the embedding sort, but perhaps the argument in favor of introspective attention of this sort is a theoretical one: one must postulate such a capacity to explain the phenomenon of introspective judgment that is epistemically privileged. This would be a theoretical posit about the existence of some capacity where positing its existence best explains the data. More discussion of this point is needed. In many ways, philosophy of mind began with Descartes' *cogito*: I think therefore I am. A natural reading of the *cogito* is that it assumes that subjects can be aware that they think via introspection. The previous discussion has, I hope, brought out how crucial attention is to introspection and how many questions about introspective attention have yet to be even articulated, let alone addressed. There is much exciting work yet to be done.

8.10 Summary

Attention is epistemically significant. There are hints of this in the discussion of propositional and doxastic justification, especially if figuring out what is true and what is reasonable to believe is an epistemic activity. For then, the issues raised concerning agency and attention in Chapter 3 arise here as well. Attention is needed, at least for doxastic justification, because forming beliefs can be an action that the agent undertakes in figuring out what to believe. In addition, fixing thoughts about consciousness also seems to involve attention, and the challenge is to figure out what sort of attention is required. On the one hand, introspective attention might just be a special way of deploying perceptual attention, a capacity that is understood fairly well. On the other hand, introspective attention might be a new form of attention, in which case more work needs to be done to understand its properties. Introspective attention is itself epistemically significant in that our understanding of it will be tied to the specific epistemic properties of introspective beliefs, such as its putative infallibility in special circumstances and its epistemic advantages over perception in other

circumstances. These are issues that philosophers are only beginning to tackle and there is much exciting work to be done.

Suggested reading

For some recent work on the epistemic role of attention, see Roessler (2011), Campbell (2011), Smithies (2011a), Siegel (forthcoming), and Silins and Siegel (Silins and Siegel, forthcoming; Siegel and Silins, forthcoming). Gertler (2011) is a volume in this series that provides a detailed overview of self-knowledge. Gertler (2012, 2001) and Chalmers (2003) discuss attention in introspection. Harman (1990) is an important source for discussion of transparency and attention. Smithies and Stoljar (2012) collects many recent discussions on the introspection of consciousness. For a recent theory of introspection where attention is discussed, see Carruthers (2011). Schwitzgebel (2011) vividly brings out how introspection can go awry.

Notes

- 1 An important discussion that I do not, unfortunately, delve into here is Susanna Siegel's recent work, especially her (Siegel, forthcoming). In this discussion, Siegel explores how attention (or what she speaks of as selection effects) can influence the rational role of experience. We will discuss her joint work with Nico Silins in a later section.
- 2 "Appropriately" here is due to cases of "bad basing," as to be discussed below.
- 3 The idea of conceptual contents in perception is rather vexing for some, who purport not to understand it. Certainly, there is some confusion about how to best formulate the issues (see Speaks 2005 for a discussion of some of the complexities). Many philosophers take perception to have nonconceptual content, even if it might also have conceptual content. For the notion of nonconceptual content, see the essays in Gunther 2003. Christopher Peacocke's (1992) notion of scenario content is a well-known account of nonconceptual content in perception.
- 4 I'm grateful to Declan Smithies for suggesting the line of argument discussed in the text. Smithies does not think that attention is necessary for propositional justification. A referee also suggests that on some conceptions of perception as having nonconceptual content, one might argue that propositions are available as reasons for belief only when nonconceptual

contents are selected for, where the role of attention is to allow for the conceptualization of those contents (see the previous footnote for the notion of nonconceptual content).

- 5 I'm grateful to Nico Silins for suggesting this line of argument.
- 6 This is not the direction that Silins and Siegel push their argument, and in personal correspondence, Silins does not object to the functional conception of attention playing the role attributed in the previous section of the text.
- 7 Silins and Siegel are targeting a plausibly different position, that doxastic justification requires attention to more than a "low degree," to finesse questions about consciousness outside of attention.
- 8 Eric Schwitzgebel (2011) has argued that perception is in fact better than introspection. He makes a strong case that introspection often goes awry and is more fallible than most philosophers allow. Sebastian Watzl and Wayne Wu (2012) have argued that many of the cases of fallibility Schwitzgebel rightfully highlights are due to inattentiveness in introspection. For example, some people will claim that the visual field is of uniform acuity, and thus that visual experience is "clear" across the visual field. If this judgment is based on introspection, then introspection gets it wrong. Of course, a moment's attention suffices to demonstrate to any subject that the visual field is not of uniform acuity, but that the periphery of the field is "blurry". Just fixate on an object and then covertly attend to the periphery.
- 9 Which is not to say that there isn't a relationalist response. Typically, the challenge just noted is to the relationalist account of phenomenal character and their typically disjunctive conception of perceptual experience. One option that some are drawn to is that in hallucination, one is aware of *uninstantiated* properties (Johnston, 2004). This is a difficult if bizarrely attractive view (at least to the author), and it would be worth asking how a Transparency Modeler might accommodate this view. We shall not do so here, but leave it as an open question.
- 10 One might wonder whether this really raises a problem for the Transparency Model since, in attending to features, we attend to properties which are universals (I owe this question to a referee for the book; see also Johnston 2004). Still, we would have to distinguish between the veridical case where we are aware of instantiated properties versus the hallucination case where we are aware of uninstantiated properties. The idea of attention to uninstantiated properties does seem to be a troubling shift from attention to instantiated properties.
- 11 For a more recent statement, see Gertler's (2012) as well as her overview on self-knowledge (2011) in the New Problems of Philosophy series.

Gertler speaks of both accounts as *acquaintance* theories, evoking Bertrand Russell's idea of knowledge by acquaintance (Russell, 1910).

- 12 Slightly adapted (replacing Gertler's "content" with "feature"). This is in fact the *preliminary* version of PSI, and Gertler later adds a fourth clause. As that clause raises complications that are not necessary for our discussion, I will omit considering it here.
- 13 Chalmers's deploys a two-dimensional semantics to supplement the analysis of a family of phenomenal concepts. These are interesting and difficult issues that I shall set aside.
- 14 Chalmers's does not characterize the phenomenal concept made available by introspective attention as a *demonstrative* concept. He agrees that attention to phenomenal features can make available demonstrative phenomenal concepts, but that these are distinct from direct phenomenal concepts. In part, this distinction is due to differing semantic analyses of the concepts, but also due to a test for cognitive significance (for discussion of the semantic analysis that appeals to a two-dimensional framework, see Chalmers 2009, chap.8). We can understand the test roughly as follows: if the identity claim X = Y is not trivial, then it is cognitively significant, and if it is cognitively significant, then "X" and "Y" are different concepts. Take now a phenomenal demonstrative concept for phenomenal feature red, this_R, and a direct phenomenal concept for the same feature, R. Chalmers claims that $this_R = this_R$ and R = R are not cognitively significant (unsurprisingly), but $this_R = R$ is. So we have two distinct concepts where R is not a demonstrative concept. In the end, the issue between Chalmers and Gertler on this point seems terminological. It is likely that Chalmers' direct phenomenal concept aligns perfectly with Gertler's pure demonstrative concept. It is not clear to me, however, that $this_R$ = $this_R$, R = R, and $this_R = R$ differ. Presumably, I have both concepts, but I cannot seem to distinguish them in thought. Try it for yourself: attend to the phenomenal character associated with red, and first deploy a demonstrative concept and then the direct concept. If there are distinct concepts, you should be able to toggle between them. An open question to readers is to see if one can have these proposed distinct thoughts.
- 15 That said, theorists who endorse relationalism about our perception of the external world might understand attention more along the lines of Chalmers and Gertler. In conversation, Chalmers indicates that he takes perceptual attention to also embed.

CONCLUSION WHAT ATTENTION IS AND WHY IT IS CENTRAL

It is perhaps appropriate here to return once more to James and his answer to the metaphysical question:

[Attention] 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, and is a condition which has a real opposite in the confused, dazed, scatterbrained state which in French is called distraction, and Zerstreutheit in German.

(James, 1890, p.403)

Given the discussion of the past eight chapters, I think James was in many ways right! Recall the five basic questions regarding attention. Here are brief responses to each in light of the discussion of this book. To underscore certain themes, I shall state the claims more boldly than perhaps the evidence and arguments warrant.

Metaphysical: What is attention?

Attention is by most accounts a selective psychological capacity. As noted in Chapter 1, there are many forms of selection that do not count as

attention. A natural specification of attentional selection, however, ties it to tasks. This link to task was uncovered in the empirical sufficient condition – selection of X for task T suffices for attention to X for T – a condition that is a shared assumption in experimental practice within the science of attention. As such, the sufficient condition provides a basic starting point for the analysis of attention. It uncovers a commonality among theorists who have lamented the possibility of defining attention. To them, one can say, "a definition is (nearly) in hand in what you already assume." That is, if theorists want to leave the mosh pit of attempts to state what attention is and not simply surrender, then the best option is to start with what everyone already knows in the experimental practice of attention: attention is, at least sometimes, selection for task.

As we have noted, the task-centered conception of attention, couched in the empirical sufficient condition, constrains interpretation of data in the neuroscience of attention. So, even in the search for basic mechanisms of attention, nothing is achieved without the empirical sufficient condition, a condition that gives neuroscientists grounds for concluding that the mechanisms and circuits that they uncover are "attentional". This points to the empirical centrality of a task-centered account of attention, one that can be expanded conceptually to a selection for action account. Of course, there are alternatives such as attention as selection for consciousness or memory, but it is likely that selection for (working) memory is subsumed by the selection for action account (selection for working memory being something necessary for much selection for action), and selection for consciousness founders on unconscious attention. The claim then is that the action-centered account provides the best current answer to the metaphysical question: attention is selection for action.

Function: What role does attention play?

One of the lessons that David Marr conveyed, in his posthumously published book, *Vision*, is that to understand capacities like vision, one must understand what that capacity is for. Without such an understanding, the cognitive science of psychological capacities can get nowhere. A computational theory, as Marr put it, is a necessary foothold for discovering mechanisms for vision at different levels of analysis. The same lesson is true of attention. What, then, is the functional role of attention? As James' quote suggests, the functional role of attention centers on its selectivity, and the usual suspects emerge: attention for action, for consciousness, for memory. One can, of course, study attentional functions in a more task-dependent and fine-grained way, say, by focusing on attention in reasoning, in perception, or in imagination, but it is likely that these more fine-grained investigations will simply uncover instantiations of the more general functions noted. Still, several of these more fine-grained functional contributions of attention are of great philosophical significance: attention's role in fixing demonstrative thought, in enabling agency, in determining and affecting the character of consciousness, in making justification possible, and in fixing introspective thought. Attention is not merely pervasive. It is fundamental to central aspects of the mind. One can put matters this way: without attention, agency, justification, certain forms of external and internal thought, and certain features of consciousness would not be possible. These are strong claims, and they merit further sustained reflection.

Properties: What are characteristic features of attention?

A slew of experimental paradigms have uncovered different features of attention: its targets, its temporal profile, its duration, its processing demands, and its interaction with other systems. Some features are notable, such as how attention is affected by the nature of the task and how it responds in different ways to different types of stimuli (e.g., direct versus symbolic cues, or feature singletons in pop-out and conjunctions in visual search). The science of attention will continue to uncover interesting features of this central psychological capacity, and as philosophers continue to explore the philosophical significance of attention, they will need to keep abreast of these developments. Still, there does seem to be a way to carve attention at its joints, namely, in the divisions between top-down and bottom-up attention and between controlled and automatic attention. These types of attention seem to involve different neural realizations, but also reflect the different sources of attention: a reliance on intentions and higher-order nonperceptual states on the one hand, and a reliance on the world on the other. Attention in that way can be both active and passive.

Mechanism: How is attention implemented?

The question of implementation can be pursued at different levels of abstraction, from abstract computational descriptions to the concrete neural realizations of attention. We have discussed, among others, Broadbent's conception of attention as a filter for selecting information for further processing, Treisman's Feature Integration Theory, with its focus on attention as binding features for object representation and awareness, Desimone and Duncane's biased competition model where attention emerges from neural competition for limited resources, Rizzolatti's Premotor Theory that takes spatial attention to result from the activation of action representations for eye movement, and a plethora of effects at the level of the activity of single neurons. The challenge for these mechanistic accounts of how attention works or is implemented will be not just whether they are adequate to the phenomenon, but also how well they can be integrated with each other. A central task for cognitive science, in which philosophers will play a critical role, is not in defining attention, for we have the basis of such a definition, but in using that definition to integrate and bridge these disparate levels of analysis. Too often within cognitive science, work at different levels fails to be bridged in illuminating ways. That is hard work, and in the case of attention, work which presents interesting challenges and fertile ground for new approaches.

Consciousness: What is the relation between attention and consciousness?

Finally, attention seems to be closely tied to consciousness. James' quote suggests that attention is essentially connected to consciousness, but some forms of attention are unconscious. So, on one reading, James was wrong: consciousness is not of attention's essence. On another reading, he was right: attention has an essential role to play in allowing us to respond to the deliverances of consciousness, for attention has a necessary role to play in our capacities to respond in general. This is not to deny that attention has a phenomenal upshot, but the precise nature of the phenomenology of attention remains a difficult question. One concrete proposal is that attention is not a consciousness. Some of the phenomenology, but is a state that affects consciousness. Some of the phenomenal effects of attention are quite disparate. Whether there is something more uniform in the phenomenology of attention, some way that attention makes things phenomenally salient, is a difficult issue on which our intuitions might simply clash at rock bottom.

Attention has another potential connection to consciousness, namely, attention as the gatekeeper of consciousness. On the one hand, the basic notion of gatekeeping can be simply expressed: one is phenomenally

conscious of X only if one attends to X. On the other hand, it is not clear exactly what this thesis comes to in detail, and it is not clear that we have strong evidence in favor of it. Some of the evidence, namely, work in the realm of inattentional blindness, is not adequate to settle the issue. At the same time, the contrary thesis-that there is phenomenology outside of attention-might seem impossible to establish, since the evidence that we have for phenomenology, namely, some form of report, relies on attention. If there is to be evidence for such phenomenal overflow, it will require ingenuity to establish. Still, gatekeeper theorists have to make more concrete exactly how to understand the limits on consciousness beyond talk of capacity limits. Many years ago, Donald Broadbent drew inspiration for psychology from the precise tools that Claude Shannon provided him in information theory. Capacity limits could be precisely quantified. We need to return to that inspiration in understanding the role of attention in consciousness and how associated capacity limits provide concrete boundaries to the character of consciousness.

So, James was right. We do know what attention is and this knowledge puts us in a position to investigate and discover what attention does, what it is good for, and how it works. The science of attention has now been established for over a century. We can, I think, look forward to a healthy philosophy of attention as well, and a positive synergy between the two approaches.

APPENDIX SHANNON INFORMATION THEORY

The basic idea behind mutual information is that when a signal Y provides us with information about some random variable X, it does so by reducing our uncertainty or surprise in respect of X. This is the fundamental idea that readers should take away. Claude Shannon's achievement was to give a precise mathematical definition of information that captures this basic idea. Consider a situation where you are uncertain about how the world is because there is a set of possibilities about how it could be, each possibility associated with a probability (e.g., it could be raining or it could be snowing). If I have information about which of these possibilities is actually the case (e.g., it is raining), then I can give you this information, say, by a signal. This signal carries an amount of information equal to the amount of uncertainty it reduces given the prior uncertainty associated with the set of possibilities. More colloquially, the signal is informative when it increases certainty about X, something that might be thought of as increasing knowledge.

To convey information by reducing uncertainty, there must be some correlation between features of that type of signal and each of the relevant possibilities. This requires that the signal can, in principle, code for each of these possibilities. In more detail, let us assume that the signal can assume some range of states $y_1 \dots y_m$ where each state is correlated with a possible state of X, the item of interest: $x_1 \dots x_n$. The first intuitive point to make is
that the amount of uncertainty in X is a function of the number of possible states X can assume. The greater the possibilities, the more potential uncertainty there is about X. Thus, consider flipping a coin, where there are two possible states, and rolling a die, where there are six possible states. Since the die has a greater number of possible states, intuitively there is more uncertainty regarding states of the die than states of the coin, once we flip them.

Shannon described this uncertainty precisely in terms of what he called the *entropy*, as given in the following formula.

(1) $H(X) = -\Sigma P[x] \log_2 P[x]$

In English, the formula states that the entropy of X, H(X), is equal to the negative of the summation of the product of (a) the logarithm of the probability of each possible state x, and (b) the probability of that specific state obtaining (I shall work through a simple example in a moment). Entropy is a function of the possible states of the relevant random variable X. When the logarithm is in base 2, entropy is reported is bits (from "binary digit").

Let us work through the example of flipping a fair coin where there are two possible states, x_1 and x_2 , namely, heads or tails. When one flips the coin, the side of the coin that lands face up $(y_1 \text{ or } y_2)$ obviously indicates for that flip, which of the two possible states (heads or tails) is achieved. This case is a special limiting case since the signal is the state of interest (normally, they are distinct), but it will suffice to draw out the central themes. Let us assume the flip yields heads. At that point, the uncertainty present before the flipping is now resolved yielding certainty: one knows which state the coin is flipped into. The flipped coin carries information in that it resolves previous uncertainty. The entropy for the coin flip turns out to be exactly 1 bit.

With biased coins, say, one weighted to heads, each flip will convey less information since, in a sense, there is less uncertainty (or surprise). After all, since the coin is biased towards heads, were one to know this, one expects heads more often than tails. The more biased the coin, the less uncertainty with the limiting case being where the coin is rigged to always land on heads. In this rigged case, there is no uncertainty resolved by flipping the coin since there is only one possible outcome: heads. The entropy in this case is equal to zero since the probability of heads is equal to 1 and the $\log_2 1 = 0$. On the other hand, as one increases the number of possible

states, H increases. For example, the information carried in a roll of the dice, where there are six equiprobable states, is H = 2.83 bits.

The information carried by a signal is then given as the mutual information, I:

(2)
$$I(X,Y) = H(X) - H(X|Y)$$

Read "I(X,Y)" as the (mutual) information of X given Y or, alternatively, the amount of information Y carries with respect to X. In fact, mutual information is symmetrical: I(X,Y) = I(Y,X), hence "mutual". Equation (2) states that the amount of information is just the difference between the entropy of X and the conditional entropy of X given signal Y. Intuitively, the conditional entropy is a measure of how much uncertainty there remains about X once one receives the signal Y, that is, how much Y has left "unsaid". If one describes H(X) as how much can be said about X (given uncertainty about it) and H(X|Y) as how much Y leaves unsaid about X, then the difference is just how much Y says about X, i.e., how much information Y conveys about X. While these notions are technically defined, it will be sufficient to understand equation (2) in the way just noted.

The maximum amount of information that a signal Y can carry about X is just the entropy associated with X, H(X), the amount of uncertainty about X. This can be seen in the case where the signal Y just is X itself (this is our limiting case). Intuitively, where Y is X, then Y leaves nothing unsaid about X. Mathematically, where X = Y, then the conditional entropy H(X|Y) = 0. Thus, in this case I(X,Y) = I(X,X) = H(X). That is, the mutual information that X carries about itself, i.e., its resolution of uncertainty regarding itself, is equal to its entropy. We then have the maximum amount of information conveyed by Y about X in the limiting case where Y = X. This is called self-information, I(X,X): the information of X given X.

We can also see that some signals convey no information about X, and accordingly I(X,Y) = 0. This occurs when H(X) = H(X|Y): the entropy of X is equal to the conditional entropy of X given Y. When does this hold? Recall the intuitive idea that H(X) is a measure of how much can be said while H(X|Y) is a measure of what Y leaves unsaid about X. Now, if Y is completely independent of X, then Y leaves everything unsaid about X. For example, I am told to report to you which side the coin falls on when I flip it, since you cannot see it, but rather than saying "heads" or "tails," I blurt out genuinely random words that are completely uncorrelated with the flips. It would be natural to accuse me, among other things, of providing you with no information about the coin. That is, my signal Y has nothing to do with X. It turns out that in this case H(X|Y) = H(X), so I(X,Y) = 0. The signal Y leaves all the uncertainty regarding X, so carries zero information.

"Information" can be ambiguous. On the one hand, "information" can signify the technical notion of mutual information, and here, the issues of quantity have a firm grip. On the other hand, "information" has a colloquial sense where it concerns meaning or content, as when speaking about a book conveying information regarding its subject matter. Thus, each sentence on this page (I hope!) conveys information in that sense, what we can call semantic information. Mutual information, however, is not semantic information (Warren Weaver (1949), in his accessible presentation of Shannon's theory, notes this point explicitly). To show this, it will be enough to show that mutual and semantic information can come apart. For example, two semantically distinct messages can carry the same amount of mutual information, while the same meaning in one context can carry some positive amount of information and in another, no information (this depends on the prior uncertainties). This should not be surprising: mutual information is a quantity; semantic information is not. So, the challenge in invoking "information" in theories of the mind and brain is separating mutual information from semantic information. Indeed, although Shannon's notion of information shaped early debates about attention, it seems that talk of information in many current psychological debates often emphasizes the semantic notion over Shannon's. Thus, one hears talk about representations of specific faces, rather than representations that carry a certain amount of information about a face. This is not to deny that there are links between mutual and semantic information, but a theory that invokes information needs to keep track of the categorical distinction between the two.

Once mutual information is defined, one can calculate the capacity limits of an information channel. First, think of a channel in terms of the probability P(Y|X) where the channel takes X as input and yields Y as output. Intuitively, the essence of the channel is just the likelihood of its giving Y as output when it receives X. This probability need not be one. Consider sending a signal in Morse code. When one sends a dash, there is some probability that it will come out as a dot (alas!). Next, the capacity limit of the channel is the maximum amount of information that it can carry. But how does one determine this? Think of X as the element that can be controlled in respect of the channel, i.e., its input. For example, if the relevant inputs are "0"s and "1"s, then think of the relevant control of input in terms of the probability of their occurring as inputs, say, being fed into the system. The question, then, is what is the maximum information that can be transmitted, given the probability of X, i.e., P(X). The maximum capacity of the channel is defined as the maximum mutual information it can carry given the probability of X as an input into the channel. Once the capacity limit of a channel is quantified, one can speak with mathematic precision about capacity limits. This provides the "new language" that Broadbent highlighted for psychology (Chapter 1). Indeed, if one aims to substantively invoke the idea of a capacity limit (e.g., in discussions of gatekeeping as in Chapter 6), one should aspire to bringing the precise notion of information to bear.

GLOSSARY

- Access Consciousness: A state or content is access conscious if it is poised for use to inform or control thought and action. One can also characterize access consciousness in terms of encoding (broadcasting) in the global workspace.
- Accuracy Condition: The accuracy condition for a representation identifies the necessary and sufficient condition for that representation to be accurate (veridical, true, satisfied). Thus, a belief that p is true if and only if p (e.g., p = "the sky is blue").
- Attended Intermediate Representations (AIR) Theory of Consciousness: A theory of consciousness due to Jesse Prinz. Phenomenal consciousness arises when intermediate perceptual representations are modulated by attention so as to be available to encoding in working memory.
- Attentional Capture: Often described as when a stimulus automatically draws attention in a bottom-up or stimulus-driven way.
- Automaticity: Automaticity is contrasted with control. In this book, automaticity is defined relative to some feature F exemplified by a process (e.g., the duration of the process), where the process is automatic relative to F if and only if it is not controlled relative to F. Attention is automatic when some of its features are.
- *Biased Competition:* Proposed neural mechanism for attention where multiple stimuli in a neuron's receptive field compete for limited processing resources such that biasing one of the stimuli, say, by top-down control, allows it to win the competition and be further processed.

- *Blindsight*: A visual condition due to damage in primary visual cortex. Blindsight subjects claim to be blind to regions of the visual field, but when prompted to "guess" about objects in their blind field, give better than chance reports. This ability is mediated by subcortical pathways that provide input to later cortical visual areas.
- *Bottom-up Attention:* Often spoken of as *endogenous* or *stimulus*-driven attention, bottom-up attention can be roughly understood as attention which is engaged due purely to sensory input. In this book, it is defined as attention whose occurrence does not depend on non-perceptual representations such as the subject's goals (e.g., an intention to attend to an object).
- *Capacity Limitations:* A fundamental constraint on processing thought to necessitate attention, namely, a limit on how much information can be processed at a time. A challenge is to provide a precise quantification of capacity limitations.
- *Capture of Attention:* A generic way to describe common situations where attention is pulled to a stimulus such as a sudden noise. Such attention is bottom-up attention.
- *Change Blindness:* An inability to detect changes. Experimentally demonstrated using displays, typically two visual images separated by a mask. There is often a substantial difference between the two images. Subjects typically do not locate changes immediately.
- *Cognitive Unison:* A theory of attention due to Christopher Mole that characterizes attention in terms of doing something *attentively*. Some task *T* is done attentively when cognitive resources are deployed in unison to serve that task.
- *Conjunction search:* A form of visual search involving identification of a target based on two features (e.g., shape and color). These features are separately exemplified by distractor objects making visual search difficult.
- *Content:* In this work, "content" often refers to representational content understood as the element of representational states that determines an accuracy condition. Sometimes, "perceptual content" is used in a general way to refer to what is perceived.
- *Content Realization Principle (CRP):* A representationalist thesis held by many scientists of consciousness: conscious content supervenes on neural information.
- *Contrast Gain:* Shift in the response of a neuron to a contrast stimulus so that the neuron is more responsive to lower contrast.
- *Control*: Opposite of automatic. Defined here in terms of features of a process: roughly, a process is controlled in respect of one of its features *F* if the process's having *F* is due to the agent's intention.
- *Coupling*: The process whereby an input is used to inform the production of a response involving parameter specification. The input is tied to a personal-level state such as the subject's visual experience.

- *Crowding:* A phenomenon where closely grouped figures result in a disruption of object perception. For example, a sequence of letters, **XAX**, is such that the flanking **X**s crowd the **A** making it difficult to discern.
- Demonstrative Concepts: Typically linguistically expressed with demonstratives such as "this" and "that". Perception-based demonstrative concepts are concepts of some X where these concepts are made available by perceptual attention to X.
- Dichotic Listening Paradigm: A paradigm pioneered by Colin Cherry where experimental subjects are required to track, i.e., attend to, one of two verbal streams presented to the ears (each ear constituting a channel). Subjects are typically required to verbally shadow one of the two streams, i.e., to verbally repeat what is said. A standard result is that few of the properties of the unattended channel are accurately reported.
- *Direct Cue*: A cue that occurs at or near the target location. Contrasted with an indirect cue.
- *Divisive Normalization:* A proposed computation performed by neurons where a neuron's response to a stimulus is divided by the response of a population of neurons that constitute its normalization pool. This computation is used to explain effects of attention on neural response.
- *Doxastic Justification:* A belief is said to be doxastically justified when it is appropriately based on propositional justification.
- *Early Selection:* The operation of selective attention early in perceptual processing of basic stimulus features and prior to higher-level processing such as semantics.
- *Early versus Late Selection:* An early debate in modern attention research concerning whether attentional filtering of information occurs early or late in perceptual processing.
- *Empirical Sufficient Condition:* If a subject *S* perceptually selects *X* for some task *T*, then *S* perceptually attends to *X* for *T*. This condition is discussed in Chapter 1.
- *Feature Attention:* Attention that is directed at perceptible properties of an object such as color or shape in vision or pitch or semantics in audition.
- *Feature Integration Theory (FIT)*: A theory due to Anne Treisman that postulates a role for attention in binding features from different visual feature maps to form a visual representation of an object.
- *Feature Search:* A form of visual search where a target is defined by a single feature.
- *Filter, Attentional:* An early conception of attention due to Donald Broadbent that emphasized attention as a filter that sifts through information, allowing only a subset for further processing.
- Fingers of Instantiation (FINST): basic mechanism postulated by Zenon Pylyshyn that connects the mind to visual objects and which makes possible representation of objects.

Firing rate: The number of spikes generated by a neuron per unit time.

Fixation Point: A location where one fixes the eyes, allowing the location to stimulate the fovea.

Fovea: A small region in the retina that provides for highest spatial acuity.

- *Gabor Patch:* A visual stimulus generated by multiplying a sinusoidal wave (function) with a Gaussian function, resulting in a contrast gradient, where the contrast in the wave decreases as it moves away from the center.
- *Gatekeeping View:* A view that characterizes attention as a gatekeeper for phenomenal consciousness. Typically the view holds that a subject is conscious of X only if the subject attends to X.
- *Global Workspace Theory*: A theory that identifies a global workspace for broadcasting signals to other systems as a necessary condition for phenomenal consciousness. The workspace is sometimes identified with working memory or with systems that input into working memory.
- *Inattentional Agnosia:* The failure of forming higher-level perceptual representations, e.g., of category, due to inattention.
- *Inattentional Amnesia:* Failure to encode information in working memory due to inattention.
- *Inattentional Apraxia:* Failure to act on information due to not attending to that information.
- *Inattentional Blindness*: Failure to be visually (phenomenally) aware of a stimulus due to inattention.
- *Indirect Cue:* Something, such as an arrow, used to indicate the location of a target. Sometimes called a symbolic cue.
- *Inhibition of Return:* A measured delay in response to stimuli at a location to which attention was previously directed. The idea is that attention is inhibited from returning to recently attended locations.
- *Information, Mutual:* A mathematically defined notion of information provided by Claude Shannon. Fundamentally, information is associated with uncertainty, with information correlated with reduction of uncertainty (for an overview, see Appendix).
- *Invalid Cue:* A cue that does not correctly indicate target location (e.g., in spatial cueing tasks).
- *Landolt Square:* A square that has a gap on one side. Used in tests for spatial acuity. In European countries, the Landolt C (a **C** oriented in various ways) is a standard test for visual acuity.
- *Late Selection:* Theories of attentional selection that hold that perceptual processing is not limited in capacity, so selection of perceptual information is post-perceptual and, hence, late.
- *Load Theory:* Theory due to Nilli Lavie that explains early and late selection effects as a function of the perceptual load in a task, i.e., the demands a task makes on processing.

- Many-Many Problem: Given that typically, agents are faced with multiple inputs and multiple possible responses, the Problem concerns how to selectively couple a specific set of inputs to a specific set of responses.
- Multiple Object Tracking (MOT): Experimental paradigm developed by Zenon Pylyshyn to test his FINST theory. Subjects track a subset of visual objects that move on a screen, typically four to five objects.
- Necessary Condition: Expressed as a conditional, "if p then q", q is a necessary condition for p. If q fails to obtain, then p does not obtain.
- *Object Attention*: Attention that is directed at objects, e.g., a sound (conversation) in audition or a material object in vision.
- *Parameter Specification:* An aspect of action where information needed to inform a response engages production mechanisms that produce that response.
- *Partial Report Paradigm:* As used by George Sperling, subjects are prompted to report only a portion of a stimulus array to which they attend. Contrasted with *total report* where subjects attempt to report the entirety of the stimulus array.
- *Perceptual Load:* The amount of available perceptual processing resources used by current processing.
- Phenomenal Consciousness: A state is phenomenally conscious if there is something that it is like to be in that state.
- *Phenomenal Character:* A feature of a conscious experience that accounts for a distinctive part of what it is like to have that experience.
- *Pop-Out*: Experimentally, tied to visual search paradigms where search is not affected by distractor set-size. This should be distinguished from colloquial talk of pop-out as attentional capture.
- Preferred Stimulus: The stimulus that drives the strongest response from a neuron when that stimulus is placed in the neuron's receptive field.
- *Premotor Theory:* A theory of spatial visual attention that posits spatial attention as a result of motor programming of eye movements to a location. An influential version is due to Giacomo Rizzolatti.
- *Reaction Time*: A standard measure in behavioral tasks, namely, how long it takes for a subject to perform a task, e.g., locating a target.
- *Receptive Field, Spatial:* In vision, the spatial receptive field of a neuron is that region of the retina, stimulation of which drives a neuron to generate spikes. Alternatively, the receptive field is that area of external space where the presence of a stimulus causes the neuron to spike.
- *Receptive Field Remapping:* In some neurons, when multiple stimuli are present in the receptive field, attention to one of the stimuli seems to cause the receptive field to contract and hence remap around the attended stimulus.
- *Relationalism:* A theory of perceptual consciousness that takes perception to be in part a relation between subject and the object of perceptual

awareness. The phenomenal character of perceptual experience is explained by a feature of the object of awareness.

- *Representationalism:* An account of phenomenal consciousness that identifies or ties the phenomenal character of a state to its representational content. Some versions identify phenomenal character with content; others take phenomenal character to merely supervene on content.
- *Response Accuracy:* A standard measure in behavioral tasks (such as reaction time), measuring the accuracy of response for a given task such as correct reports of the features of a stimulus.
- Salience: definitionally, that which attracts attention.
- Selection for Action: A theory of attention proposed by Alan Allport and Odmar Neumann that treats attention to X as the selection of X for action.
- Spatial Attention: Attention directed to a spatial location.
- Spatial Cueing: An experimental paradigm devised by Michael Posner that directs spatial attention using different types of cues (direct or indirect) that indicate the location of a target. Where attention is cued in this way with a valid cue (direct or indirect), detection of the target is improved as measured by decrease in reaction time and increase in report accuracy.
- *Spatial Frequency:* For simple visual patterns that can be represented as a sinusoidal function, the spatial frequency of the pattern is given in cycles per visual degree.
- Spike: The electrical discharge of a neuron.
- *Spotlight:* A pervasive metaphor for spatial attention, implying changes in perception in light of attention akin to spotlighting attended areas (and items) in space.
- Sufficient Condition: A sufficient condition is typically expressed as a conditional. "If p, then q" expresses the claim that p is sufficient for q. Thus, if p obtains, then q obtains.
- Supervenience: A necessary correlation that holds between X and Y where X supervenes on Y, implying that whenever there is a change in X, then there is a change in its supervenience base Y. For example, if perceptual consciousness supervenes on perceptual content, then any change in consciousness entails a change in content.
- *Top-down Attention*: Often called endogenous or goal-directed attention, in this work, it is defined as attention whose occurrence depends on a non-perceptual state, such as an intention to attend to a specific object. Its opposite is bottom-up attention.
- *Valid Cue*: A cue that correctly indicates target location. Invalid cues do not correctly indicate target location.
- *Visual Degree*: A unit of measurement of the visual field. The width of your thumb held at arms length subtends an angle of roughly two visual degrees.

286 GLOSSARY

- *Vigilance:* Often thought to be a non-selective form of attention, colloquially, a sense of readiness or openness. Technically, explication of vigilance often points to selection.
- *Visual Acuity:* A measure of spatial resolution in vision, typically tested in optometry offices using Snellen diagrams (sequences of letters of decreasing size) or Landolt figures.
- *Visual Search*: A paradigm used where a target is to be located among a set of distractors.
- *Working Memory:* A type of short-term memory that can be understood as memory for work, i.e., memory that serves ongoing behavior.

BIBLIOGRAPHY

- Abrams, Jared, Antoine Barbot, and Marisa Carrasco. 2010. "Voluntary Attention Increases Perceived Spatial Frequency." *Attention, Perception & Psychophysics* 72 (6): 1510–21. doi:10.3758/APP.72.6.1510.
- Allport, A. 1987. "Selection for Action: Some Behavioral and Neurophysiological Considerations of Attention and Action." In *Perspectives on Perception and Action*, 395–419. Hillsdale, N.J.: Lawrence Erlbaum Associates, Publishers.
- . 1993. "Attention and Control: Have We Been Asking the Wrong Questions? A Critical Review of Twenty-five Years." In Attention and Performance XIV (Silver Jubilee Volume): Synergies in Experimental Psychology, Artificial Intelligence, and Cognitive Neuroscience, 183–218. Cambridge, MA: MIT Press.
- Anderson, Britt. 2011. "There Is No Such Thing as Attention." *Frontiers in Theoretical and Philosophical Psychology* 2: 246: 1–8. doi:10.3389/ fpsyg.2011.00246.
- Anscombe, G.E.M. 1957. Intention. Oxford: Blackwell Publishers.
- Anton-Erxleben, Katharina, Jared Abrams, and Marisa Carrasco. 2010. "Evaluating Comparative and Equality Judgments in Contrast Perception: Attention Alters Appearance." *Journal of Vision* 10 (11): 1–22. doi:10.1167/10.11.6.
- 2011. "Equality Judgments Cannot Distinguish Between Attention Effects on Appearance and Criterion: a Reply to Schneider (2011)." Journal of Vision 11 (13): 1–8. doi:10.1167/11.13.8.

- Anton-Erxleben, Katharina, and Marisa Carrasco. 2013. "Attentional Enhancement of Spatial Resolution: Linking Behavioural and Neurophysiological Evidence." *Nature Reviews. Neuroscience* 14 (3): 188–200. doi:10.1038/nrn3443.
- Anton-Erxleben, Katharina, Christian Henrich, and Stefan Treue. 2007. "Attention Changes Perceived Size of Moving Visual Patterns." *Journal of Vision* 7 (11): 5.1–9. doi:10.1167/7.11.5.
- Armstrong, Katherine. 2011. "Covert Spatial Attention and Saccade Planning." In *Attention: Philosophical and Psychological Essays*, edited by Christopher Mole, Declan Smithies, and Wayne Wu, 78–96. Oxford: Oxford University Press.
- Awh, E., E. K. Vogel, and S. H. Oh. 2006. "Interactions Between Attention and Working Memory." *Neuroscience* 139 (1): 201–8.
- Azzopardi, Paul, and Alan Cowey. 1997. "Is Blindsight Like Normal, Nearthreshold Vision?" *Proceedings of the National Academy of Sciences* 94 (25): 14190–94.
- Baars, Bernard J. 1988. A Cognitive Theory of Consciousness. Cambridge: Cambridge University Press.
- -----. 1997. In the Theater of Consciousness: The Workspace of the Mind. Oxford: Oxford University Press.
- Baddeley, Alan. 2012. "Working Memory: Theories, Models, and Controversies." *Annual Review of Psychology* 63: 1–29. doi:10.1146/annurevpsych-120710-100422.
- Baddeley, Alan, and G. J. Hitch. 1974. "Working Memory." In *The Psychology of Learning and Motivation*, edited by G. H. Bower, 8:47–89. New York: Academic Press.
- Bartolomeo, Paolo, Michel Thiebaut de Schotten, and Fabrizio Doricchi. 2007. "Left Unilateral Neglect as a Disconnection Syndrome." *Cerebral Cortex* 17 (11): 2479–90. doi:10.1093/cercor/bhl181.
- Baylis, G. C., and J. Driver. 1993. "Visual Attention and Objects: Evidence for Hierarchical Coding of Location." Journal of Experimental Psychology. Human Perception and Performance 19 (3): 451–70.
- Baylis, Gordon C. 1994. "Visual Attention and Objects: Two-object Cost with Equal Convexity." *Journal of Experimental Psychology: Human Perception and Performance* 20 (1): 208–212.
- Beck, Jacob. 2012. "The Generality Constraint and the Structure of Thought." Mind 121 (483): 563–600.
- Behrmann, Marlene, Joy J. Geng, and Sarah Shomstein. 2004. "Parietal Cortex and Attention." *Current Opinion in Neurobiology* 14 (2): 212–17. doi:10.1016/j.conb.2004.03.012.

- Blaser, E., Zenon W. Pylyshyn, and A. O. Holcombe. 2000. "Tracking an Object through Feature Space." *Nature* 408 (6809): 196–99. doi:10.1038/35041567.
- Block, Ned. 1995. "On a confusion about a function of consciousness." Behavioral and Brain Sciences 18: 227-227.
- 2007a. "Consciousness, Accessibility, and the Mesh Between Psychology and Neuroscience." *Behavioral and Brain Sciences* 30 (5–6): 481–99; discussion 499–548. doi:10.1017/S0140525X07002786.
- 2007b. Consciousness, Function, and Representation. Cambridge, MA: MIT Press.
- ——. 2008. "Consciousness and Cognitive Access." Proceedings of the Aristotelian Society 108: 289–317.
- ——. 2010. "Attention and Mental Paint." Philosophical Issues 20 (1): 23–63.
- ——. 2011. "Perceptual Consciousness Overflows Cognitive Access." Trends in Cognitive Sciences 15 (12): 567–75.
- -----. 2012. "The Grain of Vision and the Grain of Attention." Thought: A Journal of Philosophy 1 (3): 170-84. doi:10.1002/tht3.28.
- ——. 2013. "Seeing and Windows of Integration." Thought: A Journal of Philosophy 2(1): 29–39. doi:10.1002/tht.62.
- Bommarito, Nicolas. 2013. "Modesty as a Virtue of Attention." *Philosophical Review* 122 (1): 93–117. doi:10.1215/00318108–1728723.
- Boynton, Geoffrey M. 2005. "Attention and Visual Perception." Current Opinion in Neurobiology 15 (4): 465–69. doi:10.1016/j.conb.2005.06.009.
- . 2009. "A Framework for Describing the Effects of Attention on Visual Responses." Vision Research 49 (10): 1129–43. doi:10.1016/j.visres.2008. 11.001.
- Brewer, B. 1999. Perception and Reason. Oxford: Oxford University Press.
- Briscoe, Robert. 2009. "Egocentric Spatial Representation in Action and Perception." Philosophy and Phenomenological Research 79 (2): 423–60.
- Brogaard, Berit. 2012. "Vision for Action and the Contents of Perception." Journal of Philosophy 109 (10): 569–87.
- Bundesen, Claus, and Thomas Habekost. 2008. Principles of Visual Attention: Linking Mind and Brain. New York: Oxford University Press.
- Burge, T. 2007. "Psychology Supports Independence of Phenomenal Consciousness." *Behavioral and Brain Sciences* 30 (5–6): 500–501.
- Burnham, Bryan R. 2007. "Displaywide Visual Features Associated with a Search Display's Appearance Can Mediate Attentional Capture." *Psychonomic Bulletin & Review* 14 (3): 392–422.

- Byrne, A. 2001. "Intentionalism Defended." *Philosophical Review* 110 (2): 199–240.
- Campbell, John. 2002. *Reference and Consciousness*. Oxford: Oxford University Press.
- —. 2003. "The Role of Demonstratives in Action Explanation." In Agency and Self-Awareness, edited by J. Roessler and N. Eilan, 150–64. Oxford: Oxford University Press.
- 2011. "Visual Attention and the Epistemic Role of Consciousness." In Attention: Philosophical and Psychological Essays, edited by Christopher Mole, Declan Smithies, and Wayne Wu, 323–42. Oxford University Press.
- Carandini, Matteo, and David J. Heeger. 2012. "Normalization as a Canonical Neural Computation." *Nature Reviews Neuroscience* 13 (1): 51–62. doi:10.1038/nrn3136.
- Carrasco, Marisa. 2006. "Covert Attention Increases Contrast Sensitivity: Psychophysical, Neurophysiological and Neuroimaging Studies." *Progress in Brain Research* 154: 33–70.
- -----. 2011. "Visual Attention: The Past 25 Years." Vision Research 51 (13): 1484-1525. doi:10.1016/j.visres.2011.04.012.
- Carrasco, Marisa, Stuart Fuller, and Sam Ling. 2008. "Transient Attention Does Increase Perceived Contrast of Suprathreshold Stimuli: a Reply to Prinzmetal, Long, and Leonhardt (2008)." *Perception & Psychophysics* 70 (7): 1151–64.
- Carrasco, Marisa, S. Ling, and S. Read. 2004. "Attention Alters Appearance." Nature Neuroscience 7 (3): 308–13.
- Carrasco, Marisa, C. Penpeci-Talgar, and M. Eckstein. 2000. "Spatial Covert Attention Increases Contrast Sensitivity Across the CSF: Support for Signal Enhancement." *Vision Research* 40 (10–12): 1203–15.
- Carrasco, Marisa, Patrick E. Williams, and Yaffa Yeshurun. 2002. "Covert Attention Increases Spatial Resolution with or Without Masks: Support for Signal Enhancement." *Journal of Vision* 2 (6): 467–79. doi:10:1167/ 2.6.4.
- Carruthers, Peter. 2011. The Opacity of Mind: An Integrative Theory of Self-Knowledge. Oxford: Oxford University Press.
- Cartwright-Finch, Ula, and Nilli Lavie. 2007. "The Role of Perceptual Load in Inattentional Blindness." *Cognition* 102 (3): 321–40. doi:10.1016/j. cognition.2006.01.002.
- Chalmers, David J. 2003. "The Content and Epistemology of Phenomenal Belief." In *Consciousness: New Philosophical Perspectives*, edited by Quentin Smith and Aleksandar Jokic. Oxford: Oxford University Press.

- ——. 2004. "The Representational Character of Experience." In The Future for Philosophy, 153–81. Oxford: Oxford University Press.
- ——. 2009. The Character of Consciousness. New York: Oxford University Press.
- Chen, Zhe. 2012. "Object-based Attention: a Tutorial Review." Attention, Perception & Psychophysics 74 (5) (July): 784–802. doi:10.3758/s13414-012-0322-z.
- Chen, Zhe, and Kyle R. Cave. 2008. "Object-based Attention with Endogenous Cuing and Positional Certainty." *Perception & Psychophysics* 70 (8): 1435– 43. doi:10.3758/PP.70.8.1435.
- Cherry, E. C. 1953. "Some Experiments on the Recognition of Speech, with One and with Two Ears." *Journal of the Acoustical Society of America* 25 (5): 975–79.
- Chun, M. M., J. D. Golomb, and N. B. Turk-Browne. 2011. "A Taxonomy of External and Internal Attention." *Annual Review of Psychology* 62: 73–101.
- Chun, M. M., and M. C. Potter. 1995. "A Two-stage Model for Multiple Target Detection in Rapid Serial Visual Presentation." *Journal of Experimental Psychology: Human Perception and Performance* 21 (1): 109–127.
- Clark, A. 2001. "Visual Experience and Motor Action: Are the Bonds Too Tight?" *Philosophical Review* 110 (4): 495–520.
- Cohen, Michael A., George A. Alvarez, and Ken Nakayama. 2011. "Naturalscene Perception Requires Attention." *Psychological Science* 22 (9): 1165–72. doi:10.1177/0956797611419168.
- Cohen, Michael A., Patrick Cavanagh, Marvin M. Chun, and Ken Nakayama. 2012. "The Attentional Requirements of Consciousness." *Trends in Cognitive Sciences* 16 (8): 411–17. doi:10.1016/j.tics.2012.06.013.
- Coltheart, M. 1980. "Iconic Memory and Visible Persistence." *Perception & Psychophysics* 27 (3): 183–228.
- Corbetta, M., and G.L. Shulman. 2002. "Control of Goal-directed and Stimulusdriven Attention in the Brain." *Nature Reviews: Neuroscience* 3: 201–205.
- Cowan, Nelson. 1995. Attention and Memory: An Integrated Framework. Oxford: Oxford University Press.
- Cumming, Sam. Manuscript "The Attentional Foundations of Coherence."
- Datta, Ritobrato, and Edgar A. DeYoe. 2009. "I Know Where You Are Secretly Attending! The Topography of Human Visual Attention Revealed with fMRI." Vision Research 49 (10): 1037–44. doi:10.1016/j.visres.2009. 01.014.
- Davidson, D. 1980. "Agency." In Essays on Actions and Events, 43–62. Oxford: Oxford University Press.

- De Brigard, Felipe. 2012. "The Role of Attention in Conscious Recollection." *Frontiers in Psychology* 3. doi:10.3389/fpsyg.2012.00029.
- De Brigard, Felipe, and Jesse Prinz. 2010. "Attention and Consciousness." Wiley Interdisciplinary Reviews 1 (1): 51–59.
- De Gardelle, Vincent, Jérôme Sackur, and Sid Kouider. 2009. "Perceptual Illusions in Brief Visual Presentations." *Consciousness and Cognition* 18 (3) (September): 569–77. doi:10.1016/j.concog.2009.03.002.
- Dehaene, Stanislas, Jean-Pierre Changeux, Lionel Naccache, Jérôme Sackur, and Claire Sergent. 2006. "Conscious, Preconscious, and Subliminal Processing: A Testable Taxonomy." *Trends in Cognitive Sciences* 10 (5): 204–11.
- Dehaene, Stanislas, M. Kerszberg, and J. P. Changeux. 1998. "A Neuronal Model of a Global Workspace in Effortful Cognitive Tasks." Proceedings of the National Academy of Sciences of the United States of America 95 (24): 14529-34.
- Dehaene, Stanislas, and Lionel Naccache. 2001. "Towards a Cognitive Neuroscience of Consciousness: Basic Evidence and a Workspace Framework." *Cognition* 79 (1): 1–37.
- Desimone, R., and J. Duncan. 1995. "Neural Mechanisms of Selective Visual Attention." *Annual Review of Neuroscience* 18: 193–222.
- Deutsch, J. A., and D. Deutsch. 1963. "Attention: Some Theoretical Considerations." *Psychological Review* 70 (1): 80–90.
- Di Lollo, Vincent. 2012. "The Feature-binding Problem Is an Ill-posed Problem." *Trends in Cognitive Sciences* 16 (6): 317–21. doi:10.1016/j. tics.2012.04.007.
- Dickie, Imogen. 2011. "Visual Attention Fixes Demonstrative Reference by Eliminating Referential Luck." In *Attention: Philosophical and Psychological Essays*, edited by Christopher Mole, Declan Smithies, and Wayne Wu, 292–322. New York: Oxford University Press.
- Dretske, Fred. 1995. Naturalizing the Mind. Cambridge, MA: MIT Press.
- ——. 2007. "What Change Blindness Teaches About Consciousness." Philosophical Perspectives 21 (1): 215–20.
- Driver, Jon. 2001. "A Selective Review of Selective Attention Research from the Past Century." *British Journal of Psychology* 92 Part 1 (February): 53-78.
- Duncan, John. 1984. "Selective Attention and the Organization of Visual Information." *Journal of Experimental Psychology. General* 113 (4): 501–17.
- Dux, P. E., and R. Marois. 2009. "The Attentional Blink: A Review of Data and Theory." Attention, Perception, & Psychophysics 71 (8): 1683–1700.

- Egly, R., J. Driver, and R. D. Rafal. 1994. "Shifting Visual Attention Between Objects and Locations: Evidence from Normal and Parietal Lesion Subjects." *Journal of Experimental Psychology. General* 123 (2): 161–77.
- Eilan, Naomi M. 1998. "Perceptual Intentionality, Attention and Consciousness." In *Current Issues in Philosophy of Mind*, edited by Anthony O'Hear, 181–202. New York: Cambridge University Press.
- Eilan, Naomi M., Christoph Hoerl, Teresa McCormack, and Johannes Roessler, ed. 2005. Joint Attention: Communication and Other Minds: Issues in Philosophy and Psychology. Oxford: Oxford University Press.
- Eriksen, C. W., and J. D. St James. 1986. "Visual Attention Within and Around the Field of Focal Attention: a Zoom Lens Model." *Perception & Psychophysics* 40 (4): 225–40.
- Evans, G. 1982. The Varieties of Reference. Oxford: Oxford University Press.
- Farah, Martha J. 2004. Visual Agnosia. Second edition. Cambridge, MA: MIT Press.
- Fodor, J. 1991. "A Theory of Content II: The Theory." In A Theory of Content and Other Essays, 89–136. Cambridge, MA: MIT Press.
- Folk, Charles L., Roger W. Remington, and James C. Johnston. 1992. "Involuntary Covert Orienting Is Contingent on Attentional Control Settings." *Journal of Experimental Psychology Human Perception and Performance* 18: 1030.
- ——. 1993. "Contingent Attentional Capture: A Reply to Yantis (1993)." Journal of Experimental Psychology: Human Perception and Performance 19 (3): 682–85. doi:10.1037/0096–1523.19.3.682.
- Freeman, Jeremy, and Denis G. Pelli. 2007. "An Escape from Crowding." Journal of Vision 7 (2): 1–14. doi:10.1167/7.2.22.
- Fuller, Stuart, and Marisa Carrasco. 2006. "Exogenous Attention and Color Perception: Performance and Appearance of Saturation and Hue." Vision Research 46 (23): 4032–47. doi:10.1016/j.visres.2006.07.014.
- Ganson, Todd, and Ben Bronner. 2012. "Visual Prominence and Representationalism." *Philosophical Studies* 164 (2): 405–18. doi:10.1007/s11098-012-9853-3.
- Gertler, Brie. 2001. "Introspecting Phenomenal States." Philosophy and Phenomenological Research 63 (2): 305–28.
- -----. 2011. Self-Knowledge. Abingdon, UK: Routledge.
- 2012. "Renewed Acquaintance." In Introspection and Consciousness, edited by Declan Smithies and Daniel Stoljar, 93–128. New York: Oxford University Press.
- Gibson, Bradley S., and Erin M. Kelsey. 1998. "Stimulus-driven Attentional Capture Is Contingent on Attentional Set for Displaywide Visual

Features." Journal of Experimental Psychology: Human Perception and Performance 24 (3): 699–706.

- Glimcher, Paul W. 2004. Decisions, Uncertainty, and the Brain: The Science of Neuroeconomics. Cambridge, MA: MIT Press.
- Gobell, Joetta, and Marisa Carrasco. 2005. "Attention Alters the Appearance of Spatial Frequency and Gap Size." *Psychological Science* 16 (8): 644–51. doi:10.1111/j.1467–9280.2005.01588.x.
- Goodale, Melvyn A., and A. David Milner. 2004. Sight Unseen: An Exploration of Conscious and Unconscious Vision. Oxford: Oxford University Press.
- Greene, Michelle R., Tommy Liu, and Jeremy M. Wolfe. 2012. "Reconsidering Yarbus: a Failure to Predict Observers' Task from Eye Movement Patterns." *Vision Research* 62 (June 1): 1–8. doi:10.1016/j.visres.2012.03.019.
- Groos, Karl. 1898. *The Play of Animals*. Translated by Elizabeth L. Baldwin. New York: D. Appleton.
- Grush, R. 2007. "A Plug for Generic Phenomenology." *Behavioral and Brain Sciences* 30 (5–6): 504–5.
- Gunther, York H. 2003. Essays on Nonconceptual Content. Cambridge, MA: MIT Press.
- Harman, G. 1990. "The Intrinsic Quality of Experience." In Philosophical Perspectives, 4 Action Theory and Philosophy of Mind, 31-52. Atascero.
- Hatfield, G. 1998. "Attention in Early Scientific Psychology." In *Visual Attention*, edited by Richard D. Wright, 3–25. Oxford: Oxford University Press.

Helmholtz, Hermann von. 1896. Handbuch der physiologischen Optik. L. Voss.

- Histed, Mark H., Vincent Bonin, and R. Clay Reid. 2009. "Direct Activation of Sparse, Distributed Populations of Cortical Neurons by Electrical Microstimulation." *Neuron* 63 (4): 508–22. doi:10.1016/j.neuron.2009. 07.016.
- Huang, Liqiang, and Harold Pashler. 2007. "A Boolean Map Theory of Visual Attention." *Psychological Review* 114 (3): 599–631. doi:10.1037/0033– 0295X.114.3.599.
- Hutchinson, J. Benjamin, and Nicholas B. Turk-Browne. 2012. "Memoryguided Attention: Control from Multiple Memory Systems." *Trends in Cognitive Sciences* 16 (12): 576–79. doi:10.1016/j.tics.2012.10.003.
- Intriligator, J., and P. Cavanagh. 2001. "The Spatial Resolution of Visual Attention." *Cognitive Psychology* 43 (3): 171–216. doi:10.1006/cogp.2001.0755.
- Itti, Laurent, and Christof Koch. 2001. "Computational Modelling of Visual Attention." Nature Reviews. Neuroscience 2 (3): 194–203. doi:10.1038/ 35058500.
- Jackendoff, Ray. 1987. Consciousness and the Computational Mind. Cambridge, MA: Bradford Book.

- Jackson, Frank. 1982. "Epiphenomenal Qualia." *Philosophical Quarterly* 32 (April): 127–36.
- James, William. 1890. *The Principles of Psychology, Volume 1.* Boston, MA: Henry Holt and Co.
- Jennings, Carolyn Dicey. 2012. "The Subject of Attention." Synthese 189 (3): 535–54.
- Jiang, Yi, Patricia Costello, Fang Fang, Miner Huang, and Sheng He. 2006. "A Gender- and Sexual Orientation-dependent Spatial Attentional Effect of Invisible Images." *Proceedings of the National Academy of Sciences* 103 (45): 17048–52. doi:10.1073/pnas.0605678103.
- Johnston, M. 2004. "The Obscure Object of Hallucination." *Philosophical Studies* 120: 113–83.
- Jonides, John, and Steven Yantis. 1988. "Uniqueness of Abrupt Visual Onset in Capturing Attention." *Perception & Psychophysics* 43 (4): 346–54.
- Kahneman, D., and A. Treisman. 1984. "Changing Views of Attention and Automaticity." In *Varieties of Attention*, edited by R Parasuraman and D. R. Davies, 29–61. New York: Academic Press.
- Kaplan, D. 1985. "Demonstratives." In *Themes from Kaplan*, 481. Oxford: Oxford University Press.
- Kastner, S., and L. G. Ungerleider. 2001. "The Neural Basis of Biased Competition in Human Visual Cortex." *Neuropsychologia* 39 (12): 1263–76.
- Keane, Brian P., and Zenon W. Pylyshyn. 2006. "Is Motion Extrapolation Employed in Multiple Object Tracking? Tracking as a Low-level, Nonpredictive Function." *Cognitive Psychology* 52 (4): 346–68. doi:10.1016/j. cogpsych.2005.12.001.

Kelly, Sean D. 2004. "Reference and Attention: A Difficult Connection." *Philosophical Studies* 120 (1-3): 277-86.

- Kentridge, R. W. 2011. "Attention Without Awareness: A Brief Review." In Attention: Philosophical and Psychological Essays, edited by Christopher Mole, Declan Smithies, and Wayne Wu, 228–46. New York: Oxford University Press.
- Kentridge, R. W., L. H. de-Wit, and C. A. Heywood. 2008. "What Is Attended in Spatial Attention?" *Journal of Consciousness Studies* 15 (4): 105–11.
- Kentridge, R. W., C.A. Heywood, and L. Weiskrantz. 1999. "Attention Without Awareness in Blindsight." *Proceedings of the Royal Society London B* 266: 1805–11.
- Kentridge, R. W., T.C.W. Nijboer, and C.A. Heywood. 2008. "Attended but Unseen: Visual Attention Is Not Sufficient for Visual Awareness." *Neuropsychologia* 46: 864–69.

- Kimchi, Ruth. 2009. "Perceptual Organization and Visual Attention." *Progress in Brain Research* 176: 15–33.
- Kimchi, Ruth, and Mary A. Peterson. 2008. "Figure-Ground Segmentation Can Occur Without Attention." *Psychological Science* 19 (7): 660–68. doi:10.1111/j.1467–9280.2008.02140.x.
- Kind, Amy. 2003. "What's So Transparent About Transparency?." *Philosophical Studies* 115.3: 225-244.
- Klein, R. 1980. "Does Oculomotor Readiness Mediate Cognitive Control of Visual Attention." In Attention and Performance VIII, edited by R. S. Nickerson, 8:259–76. Hillsdale, N.J.: Lawrence Erlbaum Associates, Publishers.
- Klein, Raymond M., and Michael A. Lawrence. 2012. "On the Modes and Domains of Attention." In *Cognitive Neuroscience of Attention, Second Edition*, edited by Michael I. Posner and Liqiang Huang, 11–28. New York: Guilford Press.
- Koch, Christof, and Naotsugu Tsuchiya. 2007. "Attention and Consciousness: Two Distinct Brain Processes." *Trends in Cognitive Sciences* 11 (1): 16–22. doi:10.1016/j.tics.2006.10.012.
- Koralus, Philipp. 2014. "The Erotetic Theory of Attention: Questions, Focus, and Distraction." *Mind and Language* 29(1): 26–50. doi:10/1111/mila.12040.
- Kouider, Sid, Vincent de Gardelle, Jérôme Sackur, and Emmanuel Dupoux. 2010. "How Rich Is Consciousness? The Partial Awareness Hypothesis." *Trends in Cognitive Sciences* 14 (7): 301–7. doi:10.1016/j.tics.2010. 04.006.
- Koulakov, Alexei A., Brian E. Kolterman, Armen G. Enikolopov, and Dmitry Rinberg. 2011. "In Search of the Structure of Human Olfactory Space." *Frontiers in Systems Neuroscience* 5: 65. doi:10.3389/fnsys.2011.00065.
- Lamme, Victor A. F. 2004. "Separate Neural Definitions of Visual Consciousness and Visual Attention; a Case for Phenomenal Awareness." Neural Networks: The Official Journal of the International Neural Network Society 17 (5–6) (July): 861–72. doi:10.1016/j.neunet.2004.02.005.
- Lamme, Victor A. F. 2010. "How Neuroscience Will Change Our View on Consciousness." *Cognitive Neuroscience* 1 (3): 204–20. doi:10.1080/ 17588921003731586.
- Land, M. F. 2006. "Eye Movements and the Control of Actions in Everyday Life." *Progress in Retinal and Eye Research* 25: 296–324.
- Landman, Rogier, Henk Spekreijse, and Victor A. F. Lamme. 2003. "Large Capacity Storage of Integrated Objects before Change Blindness." *Vision Research* 43 (2): 149–64.

- Lavie, Nilli. 2005. "Distracted and Confused?: Selective Attention Under Load." *Trends in Cognitive Sciences* 9 (2): 75–82. doi:10.1016/j.tics. 2004.12.004.
- Lavie, Nilli, and Yehoshua Tsal. 1994. "Perceptual Load as a Major Determinant of the Locus of Selection in Visual Attention." *Perception & Psychophysics* 56 (2): 183–97. doi:10.3758/BF03213897.
- Lee, Joonyeol, and John H. R. Maunsell. 2009. "A Normalization Model of Attentional Modulation of Single Unit Responses." *PLoS ONE* 4 (2): e4651. doi:10.1371/journal.pone.0004651.
- Lerman, Hemdat. manuscript "Visual Attention and the Phenomenology of Visual Experience". University of Warwick.
- Li, Fei Fei, Rufin VanRullen, Christof Koch, and Pietro Perona. 2002. "Rapid Natural Scene Categorization in the Near Absence of Attention." Proceedings of the National Academy of Sciences of the United States of America 99 (14): 9596–9601. doi:10.1073/pnas.092277599.
- Liu, Taosheng, Jared Abrams, and Marisa Carrasco. 2009. "Voluntary Attention Enhances Contrast Appearance." *Psychological Science* 20 (3): 354–62.
- Losier, Bruno J.W., and Raymond M. Klein. 2001. "A Review of the Evidence for a Disengage Deficit Following Parietal Lobe Damage." *Neuroscience & Biobehavioral Reviews* 25 (1): 1–13. doi:10.1016/S0149-7634(00)00046-4.
- Luck, S. J., L. Chelazzi, S. A. Hillyard, and R. Desimone. 1997. "Neural Mechanisms of Spatial Selective Attention in Areas V1, V2, and V4 of Macaque Visual Cortex." *Journal of Neurophysiology* 77 (1): 24–42.
- Machamer, P., L. Darden, and C. Craver. 2000. "Thinking About Mechanisms." *Philosophy of Science* 67: 1–25.
- Mack, Arien, and Jason Clarke. 2012. "Gist Perception Requires Attention." *Visual Cognition* 20 (3): 300–327. doi:10.1080/13506285.2012.666578.
- Mack, Arien, and Irvin Rock. 1998. *Inattentional Blindness*. Cambridge, MA: MIT Press.
- Marr, D. 1982. Vision. San Francisco: W. H. Freeman and Company.
- Marshall, J. C., and P. W. Halligan. 1988. "Blindsight and Insight in Visuospatial Neglect." *Nature* 336 (6201): 766–67. doi:10.1038/336766a0.
- Martens, S., and B. Wyble. 2010. "The Attentional Blink: Past, Present, and Future of a Blind Spot in Perceptual Awareness." *Neuroscience* & *Biobehavioral Reviews* 34 (6): 947–57.
- Martin, Michael G. F. 2002. "The Transparency of Experience." Mind & Language 17.4: 376-425.
- Martínez, A., W. Teder-Sälejärvi, M. Vazquez, S. Molholm, J.J. Foxe, D. C. Javitt, F. Di Russo, M. S. Worden, and S. A. Hillyard. 2006. "Objects Are

Highlighted by Spatial Attention." Journal of Cognitive Neuroscience 18 (2): 298-310. doi:10.1162/089892906775783642.

- Matthen, Mohan. 2006. "On Visual Experience of Objects: Comments on John Campbell's Reference and Consciousness." *Philosophical Studies* 127 (2): 195–220.
- McAdams, C. J., and John H. R. Maunsell. 1999. "Effects of Attention on the Reliability of Individual Neurons in Monkey Visual Cortex." *Neuron* 23 (4) (August): 765–73.
- McDowell, J. 1994. *Mind and World*. Cambridge, MA: Harvard University Press.
- Millikan, R.G. 1984. Language, Thought and Other Biological Categories: New Foundations for Realism. Cambridge, MA: MIT Press.
- Milner, A. David, and Melvyn A. Goodale. 1995. *The Visual Brain in Action*. Oxford: Oxford University Press.
- -----. 2006. The Visual Brain in Action. 2nd ed. Oxford: Oxford University Press.
- Mishkin, M., L. G. Ungerleider, and K. A. Macko. 1983. "Object Vision and Spatial Vision: Two Cortical Pathways." *Trends in Neurosciences* 6: 414–17.
- Mole, Christopher. 2008a. "Attention and Consciousness." Journal of Consciousness Studies 15 (4): 86–104.
- 2008b. "Attention in the Absence of Consciousness?" Trends in Cognitive Sciences 12 (2): 44.
- ——. 2009. "Illusions, Demonstratives and the Zombie Action Hypothesis," Mind 118 (472): 995–1011.

——. 2011. Attention Is Cognitive Unison: An Essay in Philosophical Psychology. New York: Oxford University Press.

- 2013. "Review of Jesse J. Prinz, The Conscious Brain." Notre Dame Phillosophical Reviews.
- Molyneux, Bernard. 2009. "Why Experience Told Me Nothing about Transparency." *Nous* 43.1: 116-136.
- Montagna, Barbara, Franco Pestilli, and Marisa Carrasco. 2009. "Attention Trades Off Spatial Acuity." *Vision Research* 49 (7) (March): 735–45.
- Moore, Cathleen M., Steven Yantis, and Barry Vaughan. 1998. "Object-based Visual Selection: Evidence from Perceptual Completion." *Psychological Science* 9 (2): 104–10.
- Moore, Tirin, and Katherine M. Armstrong. 2003. "Selective Gating of Visual Signals by Microstimulation of Frontal Cortex." *Nature* 421 (6921) (January 23): 370–73. doi:10.1038/nature01341.

- Moore, Tirin, and Martha J. Fallah. 2001. "Control of Eye Movements and Spatial Attention." *Proceedings of the National Academy of Sciences of the United States of America* 98 (3) (January 30): 1273–76. doi:10.1073/ pnas.021549498.
- Moran, J., and R. Desimone. 1985. "Selective Attention Gates Visual Processing in the Extrastriate Cortex." *Science* 229 (4715) (August 23): 782–84.
- Moray, Neville. 1959. "Attention in Dichotic Listening: Affective Cues and the Influence of Instructions." *Quarterly Journal of Experimental Psychology* 11 (1): 56–60. doi:10.1080/17470215908416289.
- Munoz, Douglas P. 2002. "Commentary: Saccadic Eye Movements: Overview of Neural Circuitry." *Progress in Brain Research* 140: 89–96. doi:10.1016/ S0079-6123(02)40044-1.
- Naccache, Lionel, and Stanislas Dehaene. 2007. "Reportability and Illusions of Phenomenality in the Light of the Global Neuronal Workspace Model." *Behavioral and Brain Sciences* 30 (5–6): 518–20. doi:10.1017/S0140525 X07002993.
- Nagel, Thomas. 1974. "What Is It Like to Be a Bat?" *Philosophical Review* 83 (October): 435-50.
- Nakayama, Ken, and Gerald H. Silverman. 1986. "Serial and Parallel Processing of Visual Feature Conjunctions" 320 (6059) (March 20): 264–65. doi:10.1038/320264a0.

Nanay, Bence. 2010. "Attention and Perceptual Content." Analysis 70 (2): 263-70.

- Navon, David, and Daniel Gopher. 1979. "On the Economy of the Humanprocessing System." *Psychological Review* 86 (3): 214–55. doi:10.1037/ 0033-0295X.86.3.214.
- Neisser, Ulric. 1967. Cognitive Psychology. London: Prentice-Hall International, Inc.
- Neisser, Ulric, and Robert Becklen. 1975. "Selective Looking: Attending to Visually Specified Events." *Cognitive Psychology* 7 (4): 480–94. doi:10.1016/0010-0285(75)90019-5.
- Neta, Ram. 2010. "The Basing Relation." In *The Routledge Companion to Epistemology*, edited by Sven Bernecker and Duncan Pritchard, 109–18. Abingdon, UK: Routledge.
- Neumann, O. 1987. "Beyond Capacity: A Functional View of Attention." In *Perspectives on Perception and Action*, 361–94. Hillsdale: Lawrence Erlbaum Associates, Publishers.
- Ni, Amy M., Supratim Ray, and John H. R. Maunsell. 2012. "Tuned Normalization Explains the Size of Attention Modulations." *Neuron* 73 (4): 803– 13. doi:10.1016/j.neuron.2012.01.006.

Nickel, Bernhard. 2007. "Against Intentionalism." *Philosophical Studies* 136 (3): 279–304.

Noë, Alva. 2002. Is the Visual World a Grand Illusion? Thorverton: Imprint Academic.

Norman, Donald A. 1968. "Toward a Theory of Memory and Attention." *Psy*chological Review 75 (6): 522-36. doi:10.1037/h0026699.

- Norman, Liam J., Charles A. Heywood, and Robert W. Kentridge. 2013. "Object-Based Attention Without Awareness." *Psychological Science* 24 (6): 836–43. doi:10.1177/0956797612461449.
- Noudoost, Behrad, Mindy H. Chang, Nicholas A. Steinmetz, and Tirin Moore. 2010. "Top-down Control of Visual Attention." *Current Opinion in Neurobiology* 20 (2): 183–90. doi:10.1016/j.conb.2010.02.003.
- Oken, B. S., M. C. Salinsky, and S. M. Elsas. 2006. "Vigilance, Alertness, or Sustained Attention: Physiological Basis and Measurement." *Clinical Neurophysiology: Official Journal of the International Federation of Clinical Neurophysiology* 117 (9): 1885–1901. doi:10.1016/j.clinph.2006.01.017.
- Oliva, A. 2005. "Gist of the Scene." In *Neurobiology of Attention*, edited by Laurent Itti, Geraint Rees, and John Tsotsos, 251–56. Burlington, MA: Elsevier Academic Press.
- Overgaard, Morten. 2011. "Visual Experience and Blindsight: a Methodological Review." *Experimental Brain Research* 209 (4): 473–79. doi:10.1007/ s00221-011-2578-2.
- Parasuraman, Raja, Joel S. Warm, and Judi E. See. 1998. "Brain Systems of Vigilance." In *The Attentive Brain*, edited by Parasuraman Parasuraman, 221–56. Cambridge, MA: MIT Press.

Pashler, H.E. 1998. The Psychology of Attention. Cambridge, MA: MIT Press.

- Pautz, Adam. 2010. "An Argument for the Intentional View of Visual Experience." In *Perceiving the World*, edited by Bence Nanay. Oxford: Oxford University Press.
- Peacocke, C. 1992. A Study of Concepts. Cambridge, MA: MIT Press.
- ——. 1998. "Conscious Attitudes, Attention, and Self-Knowledge." In Knowing Our Own Minds, 63–98. Oxford: Oxford University Press.
- Pelli, Denis G., and Katharine A Tillman. 2008. "The Uncrowded Window of Object Recognition." *Nature Neuroscience* 11 (10) (October): 1129–35.
- Pestilli, Franco, and Marisa Carrasco. 2005. "Attention Enhances Contrast Sensitivity at Cued and Impairs It at Uncued Locations." *Vision Research* 45 (14): 1867–75. doi:10.1016/j.visres.2005.01.019.
- Petersen, Steven E., and Michael I. Posner. 2012. "The Attention System of the Human Brain: 20 Years After." *Annual Review of Neuroscience* 35 (1): 73–89. doi:10.1146/annurev-neuro-062111-150525.

- Petrov, Yury, and Ariella V. Popple. 2007. "Crowding Is Directed to the Fovea and Preserves Only Feature Contrast." *Journal of Vision* 7 (2): 8.1–9. doi:10.1167/7.2.8.
- Phillips, Ian. 2011a. "Perception and Iconic Memory: What Sperling Doesn't Show." *Mind and Language* 26 (4): 381–411.
- 2011b. "Attention and Iconic Memory." In Attention: Philosophical and Psychological Essays, edited by Christopher Mole, Declan Smithies, and Wayne Wu, 204–27. New York: Oxford University Press.
- -----. 2012. "Attention to the Passage of Time." *Philosophical Perspectives* 26 (1): 277-308.
- Posner, Michael I. 1980. "Orienting of Attention." The Quarterly Journal of Experimental Psychology 32 (1): 3-25.
- Posner, Michael I. 2011. Cognitive Neuroscience of Attention. 2nd edition. The Guilford Press: New York City.
- Posner, Michael I., and Steven E. Petersen. 1990. "The Attention System of the Human Brain." *Annual Review of Neuroscience* 13 (1): 25–42. doi:10.1146/annurev.ne.13.030190.000325.
- Pouget, Pierre, Iwona Stepniewska, Erin A. Crowder, Melanie W. Leslie, Erik E. Emeric, Matthew J. Nelson, and Jeffrey D. Schall. 2009. "Visual and Motor Connectivity and the Distribution of Calcium-binding Proteins in Macaque Frontal Eye Field: Implications for Saccade Target Selection." *Frontiers in Neuroanatomy* 3: 2. doi:10.3389/neuro.05. 002.2009.
- Prinz, Jesse. 2012. The Conscious Brain. Oxford: Oxford University Press.
- Prinzmetal, William, Virginia Long, and James Leonhardt. 2008. "Involuntary Attention and Brightness Contrast." *Perception & Psychophysics* 70 (7) (October): 1139–50. doi:10.3758/PP.70.7.1139.
- Pylyshyn, Zenon W. 2007. Things and Places: How the Mind Connects With the World. Cambridge, MA: MIT Press.
- Quinlan, Philip T. 2003. "Visual Feature Integration Theory: Past, Present, and Future." *Psychological Bulletin* 129 (5): 643–73. doi:10.1037/0033–2909. 129.5.643.
- Reynolds, John H., Leonardo Chelazzi, and Robert Desimone. 1999. "Competitive Mechanisms Subserve Attention in Macaque Areas V2 and V4." *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience* 19 (5): 1736–53.
- Reynolds, John H., and David J. Heeger. 2009. "The Normalization Model of Attention." *Neuron* 61 (2) (January 29): 168–85. doi:10.1016/j.neuron. 2009.01.002.

- Reynolds, John H., Tatiana Pasternak, and Robert Desimone. 2000. "Attention Increases Sensitivity of V4 Neurons." *Neuron* 26 (3): 703–14. doi:10.1016/S0896–6273(00)81206–4.
- Richards, Bradley. 2013. "Identity-Crowding and Object-Seeing: A Reply to Block." Thought: A Journal of Philosophy 2 (1): 9–19. doi:10.1002/tht3.57.
- Rizzolatti, Giacomo, Lucia Riggio, and Bori M. Sheliga. 1994. "Space and Selective Attention." In Attention and Performance XV: Conscious and Nonconscious Information Processing, edited by Carlo Umiltà and Morris Moscovitch, 231–65. Cambridge, MA: MIT Press.
- Robbins, Trevor W. 1998. "Arousal and Attention: Psychopharmacological and Neuropsychological Studies in Experimental Animals." In *The Attentive Brain*, 189–220. Cambridge, MA: The MIT Press.
- Rock, Irvin, and Daniel Gutman. 1981. "The Effect of Inattention on Form Perception." Journal of Experimental Psychology: Human Perception and Performance 7 (2): 275–85.
- Roessler, Johannes. 2011. "Perceptual Attention and the Space of Reasons." In *Attention: Philosophical and Psychological Essays*, edited by Christopher Mole, Declan Smithies, and Wayne Wu, 274–92. New York: Oxford University Press.
- Rolfs, Martin. 2009. "Microsaccades: Small Steps on a Long Way." Vision Research 49 (20) (October 15): 2415–41. doi:10.1016/j.visres.2009.08.010.
- Ruff, Christian. 2011. "A Systems-Neuroscience View of Attention." In Attention: Philosophical and Psychological Essays, edited by Christopher Mole, Declan Smithies, and Wayne Wu, 1–23. New York: Oxford University Press.
- Russell, Bertrand. 1910. "Knowledge by Acquaintance and Knowledge by Description." *Proceedings of the Aristotelian Society* 11: 108–28.
- Schacter, Daniel L. 2012. "Constructive Memory: Past and Future." *Dialogues in Clinical Neuroscience* 14 (1): 7–18.
- Schacter, Daniel L., K. A. Norman, and W. Koutstaal. 1998. "The Cognitive Neuroscience of Constructive Memory." *Annual Review of Psychology* 49: 289–318. doi:10.1146/annurev.psych.49.1.289.
- Schenk, T., and R. D. McIntosh. 2010. "Do We Have Independent Visual Streams for Perception and Action." *Cognitive Neuroscience* 1 (1): 52–78.
- Schneider, Keith A., and Marcell Komlos. 2011. "Attention Alters Decision Criteria but Not Appearance: a Reanalysis of Anton-Erxleben, Abrams, and Carrasco (2010)." *Journal of Vision* 11 (13): 1–10. doi:10.1167/11.13.7.
- Schneider, W., and R.M. Shiffrin. 1977. "Controlled and Automatic Human Information Processing: I. Detection, Search and Attention." *Psychological Review* 84 (1): 1–66.

- Scholl, Brian J. 2001. "Objects and Attention: The State of the Art." *Cognition* 80 (1–2): 1–46. doi:10.1016/S0010-0277(00)00152-59.
 - —. 2009. "What Have We Learned About Attention from Multiple Object Tracking (and Vice Versa)?" In *Computation, Cognition, and Pylyshyn*, edited by D. Dedrick and L. Trick, 49–78. Cambridge, MA: MIT Press.
- Scholl, Brian J., Z. W. Pylyshyn, and J. Feldman. 2001. "What Is a Visual Object? Evidence from Target Merging in Multiple Object Tracking." *Cognition* 80 (1–2): 159–77.
- Schwitzgebel, Eric. 2011. Perplexities of Consciousness. Cambridge, MA: MIT Press.
- Seemann, Axel. 2011. Joint Attention: New Developments in Psychology, Philosophy of Mind, and Social Neuroscience. Cambridge, MA: MIT Press.
- Shanahan, Murray, and Bernard Baars. 2007. "Global Workspace Theory Emerges Unscathed." *Behavioral and Brain Sciences* 30 (5–6): 524–25. doi:10.1017/S0140525X07003056.
- Shapiro, Lawrence A. 2000. "Multiple Realizations." *Journal of Philosophy* 97 (12): 635–54.

-----. 2004. The Mind Incarnate. Cambridge, MA: MIT Press.

- Shomstein, Sarah. 2012. "Object-based Attention: Strategy Versus Automaticity." Wiley Interdisciplinary Reviews: Cognitive Science 3 (2): 163–69.
- Shomstein, Sarah, and Steven Yantis. 2002. "Object-based Attention: Sensory Modulation or Priority Setting?" *Perception & Psychophysics* 64 (1): 41–51.
- Shulman, Gordon L., and Maurizio Corbetta. 2012. "Two Attentional Networks: Identification and Function Within a Larger Cognitive Architecture." In Cognitive Neuroscience of Attention, Second Edition, edited by Michael I. Posner and Ligiang Huang, 113–28. New York: Guilford Press.
- Siegel, Susanna. forthcoming. "Can Selection Effects on Experience Influence Its Rational Role?" In *Oxford Studies in Epistemology Volume 4*, edited by Tamar Gendler. Oxford.
- ——. 2010. The Contents of Visual Experience. New York: Oxford University Press.
- Siegel, Susanna, and Nicholas Silins. forthcoming. "Attention and Perceptual Justification." In *Festschrift for Ned Block*, edited by Adam Pautz and Daniel Stoljar. Cambridge, MA: MIT Press.
- Siewert, Charles. 2004. "Is experience transparent?" *Philosophical Studies* 117.1: 15-41.
- Silins, Nicholas, and Susanna Siegel. forthcoming. "Consciousness, Attention, and Justification." In *Contemporary Perspectives on Scepticism and Perceptual Jusification*, edited by Elia Zardini and Dylan Dodd. Oxford: Oxford University Press.

- Simons, Daniel J. 2000. "Attentional Capture and Inattentional Blindness." *Trends in Cognitive Sciences* 4 (4): 147–55. doi:10.1016/S1364–6613(00) 01455–58.
- Simons, Daniel J., and M. S. Ambinder. 2005. "Change Blindness Theory and Consequences." *Current Directions in Psychological Science* 14 (1): 44–48.
- Sligte, Ilja G., H. Steven Scholte, and Victor A. F. Lamme. 2008. "Are There Multiple Visual Short-term Memory Stores?" *PloS One* 3 (2): e1699. doi:10.1371/journal.pone.0001699.
- Smith, Daniel T., and Thomas Schenk. 2012. "The Premotor Theory of Attention: Time to Move On?" *Neuropsychologia* 50 (6): 1104–14. doi:10.1016/ j.neuropsychologia.2012.01.025.
- Smithies, Declan. 2011a. "What Is the Role of Consciousness in Demonstrative Thought?" Journal of Philosophy 108 (1): 5-34.
- 2011b. "Attention Is Rational-Access Consciousness." In Attention: Philosophical and Psychological Essays, edited by Christopher Mole, Declan Smithies, and Wayne Wu. New York: Oxford University Press.
- Smithies, Declan, and Daniel Stoljar. 2012. *Introspection and Consciousness*. New York: Oxford University Press, 247–73.
- Sosa, Ernest. 2003. "Privileged Access." In *Consciousness: New Philosophical Perspectives*, edited by Quentin Smith and Aleksandar Jokic, 238–51. New York: Oxford University Press.
- Speaks, Jeff. 2005. "Is There a Problem About Nonconceptual Content?" *Philosophical Review* 114: 359–98.
- ——. 2010. "Attention and Intentionalism." Philosophical Quarterly 60 (239): 325–42.

Spence, C. J., and J. Driver. 1994. "Covert Spatial Orienting in Audition: Exogenous and Endogenous Mechanisms." *Journal of Experimental Psychology. Human Perception and Performance* 20 (3): 555–74.

- Sperling, George. 1960. "The Information Available in Brief Visual Presentations." *Psychological Monographs: General and Applied* 74 (11): 1–29.
- Stazicker, J. 2011. "Attention, Visual Consciousness and Indeterminacy." *Mind* & Language 26 (2): 156–84.
- -----. 2011. "Attention & the Indeterminacy of Visual Experience". Ph.D., United States – California: University of California, Berkeley.
- Stoljar, Daniel. 2004. "The Argument from Diaphanousness" in New Essays in Philosophy of Language and Mind. Maite Ezcurdia, Robert J. Stainton, and Christopher Viger, eds. Canadian Journal of Philosophy, Supplement 30.
- Störmer, Viola S., John J. McDonald, and Steven A. Hillyard. 2009. "Crossmodal Cueing of Attention Alters Appearance and Early Cortical

Processing of Visual Stimuli." *Proceedings of the National Academy of Sciences* 106 (52) (December 29): 22456–61. doi:10.1073/pnas.0907573106.

Sutherland, Stuart. 1998. "Feature Selection." *Nature* 392 (6674): 350. doi:10.1038/ 32817.

- Taylor, J. H. 2013. "Is Attention Necessary and Sufficient for Phenomenal Consciousness?" *Journal of Consciousness Studies* 20(11–12): 173–94.
- -----. 2013. "Is the Grain of Vision Finer Than the Grain of Attention? Response to Block." *Thought: A Journal of Philosophy*. doi:10.1002/tht.59.
- Theeuwes, Jan. 2010. "Top-down and Bottom-up Control of Visual Selection." *Acta Psychologica* 135 (2): 77-99.
- Thompson, Kirk G., and Narcisse P. Bichot. 2005. "A Visual Salience Map in the Primate Frontal Eye Field." In *Progress in Brain Research*, edited by M. Kamermans J. van Pelt, Volume 147:249–62. Elsevier. http://www.scien cedirect.com/science/article/pii/S0079612304470198.
- Treisman, Anne M. 1960. "Contextual Cues in Selective Listening." Quarterly Journal of Experimental Psychology 12 (4): 242–48.
- . 1969. "Strategies and Models of Selective Attention." Psychological Review 76 (3): 282–99. doi:10.1037/h0027242.
- ——. 1988. "Features and Objects: The Fourteenth Bartlett Memorial Lecture." Quarterly Journal of Experimental Psychology 40: 201–37.

-----. 1996. "The Binding Problem." Current Opinion in Neurobiology 6: 171-8.

Treisman, Anne M., and G. Gelade. 1980. "A Feature-Integration Theory of Attention." *Cognitive Psychology* 12: 97–136.

Treisman, Anne M., and S. Gormican. 1988. "Feature Analysis in Early Vision: Evidence from Search Asymmetries." *Psychological Review* 95 (1): 15–48.

- Treisman, Anne M., D. Kahneman, and J. Burkell. 1983. "Perceptual Objects and the Cost of Filtering." *Perception & Psychophysics* 33 (6): 527–32.
- Tse, Peter U. 2005. "Voluntary Attention Modulates the Brightness of Overlapping Transparent Surfaces." *Vision Research* 45 (9): 1095–98.
- Tsotsos, John K. 2011. A Computational Perspective on Visual Attention. Cambridge, MA: MIT Press.
- Tsotsos, John, and Albert Rothenstein. 2011. "Computational Models of Visual Attention." *Scholarpedia* 6 (1): 6201.
- Tsuchiya, Naotsugu, Ned Block, and Christof Koch. 2012. "Top-down Attention and Consciousness: Comment on Cohen et Al." *Trends in Cognitive Sciences* (September 28) 16: 527.
- Turatto, Massimo, Massimo Vescovi, and Matteo Valsecchi. 2007. "Attention Makes Moving Objects Be Perceived to Move Faster." Vision Research 47 (2) (January): 166–78.

- Turri, John. 2010. "On the Relationship Between Propositional and Doxastic Justification." *Philosophy and Phenomenological Research* 80 (2): 312–26.
- Tye, Michael. 1995. Ten Problems of Consciousness: A Representational Theory of the Phenomenal Mind. Cambridge, MA: MIT Press.
- -----. 2000. Consciousness, Color and Content. Cambridge MA: MIT Press.
- 2003. "Blurry Images, Double Vision, and Other Oddities: New Problems for Representationalism?" In Consciousness: New Philosophical Perspectives, edited by Quentin Smith and Aleksandar Jokic, 7–32. Oxford: Oxford University Press.
- -----. 2009. "A New Look at the Speckled Hen." Analysis 69 (2): 258-63.
- Ungerleider, L. G. & Mishkin, M. 1982. "Two Cortical Visual Systems". In Analysis of Visual Behavior (ed. D.J. Ingle, M.A. Goodale & R.J.W. Mansfeld). Cambridge, MA: MIT Press 549-86.
- Van Essen, David C. 2005. "A Population-Average, Landmark- and Surfacebased (PALS) Atlas of Human Cerebral Cortex." *NeuroImage* 28 (3): 635–62.
- VanMarle, Kristy, and Brian J. Scholl. 2003. "Attentive Tracking of Objects Versus Substances." *Psychological Science* 14 (5) (September): 498–504.
- Vecera, Shaun P., and Matthew Rizzo. 2003. "Spatial Attention: Normal Processes and Their Breakdown." *Neurologic Clinics* 21 (3): 575–607.
- Wannig, Aurel, Liviu Stanisor, and Pieter R. Roelfsema. 2011. "Automatic Spread of Attentional Response Modulation Along Gestalt Criteria in Primary Visual Cortex." *Nature Neuroscience* 14 (10): 1243–44.
- Watzl, Sebastian. 2011. "Attention as Structuring of the Stream of Consciousness." In *Attention: Philosophical and Psychological Essays*, edited by Christopher Mole, Declan Smithies, and Wayne Wu, 145–73. New York: Oxford University Press.
- Watzl, Sebastian, and Wayne Wu. 2012. "Perplexities of Consciousness, by Eric Schwitzgebel." *Mind* 121 (482): 524–29.
- Weaver, W. 1949. "Recent Contributions to The Mathematical Theory of Communication." In *The Mathematical Theory of Communication*, 1–28. Urbana: University of Illinois Press.
- Weiskrantz, Lawrence. 1999. Consciousness Lost and Found: A Neuropsychological Exploration. Oxford: Oxford University Press.
- Whitney, David, and Dennis M. Levi. 2011. "Visual Crowding: a Fundamental Limit on Conscious Perception and Object Recognition." *Trends in Cognitive Sciences* 15 (4): 160–68.

- Wolfe, Jeremy M. 1999. "Inattentional Amnesia." In *Fleeting Memories: Cognition* of Brief Visual Stimuli, edited by Coltheart, Max, 71–94. Cambridge, MA: MIT Press.
- Wolfe, Jeremy M, and Todd S. Horowitz. 2004. "What Attributes Guide the Deployment of Visual Attention and How Do They Do It?" *Nature Reviews. Neuroscience* 5 (6): 495–501.
- Wolfe, Jeremy M. 1994. "Guided Search 2.0 A Revised Model of Visual Search." *Psychonomic Bulletin & Review* 1 (2): 202–38.
- 2003. "Moving Towards Solutions to Some Enduring Controversies in Visual Search." Trends in Cognitive Sciences 7 (2): 70–76.
- ——. 2012. "The Binding Problem Lives on: Comment on Di Lollo." Trends in Cognitive Sciences 16 (6): 307–8.
- Wolpert, D.M., and Z. Ghahramani. 2000. "Computational Principles of Movement Neuroscience." *Nature Neuroscience* 3 Suppl: 1212–17.
- Womelsdorf, Thilo, and Pascal Fries. 2007. "The Role of Neuronal Synchronization in Selective Attention." *Current Opinion in Neurobiology* 17 (2) (April): 154–60.
- Wright, Richard D., and Lawrence M. Ward. 2008. *Orienting of Attention*. New York: Oxford University Press US.
- Wu, Wayne. forthcoming. "Against Division: Consciousness, Information and the Visual Streams." *Mind and Language*
- . 2008. "Visual Attention, Conceptual Content, and Doing It Right." Mind
 117 (468): 1003–33.
- ——. 2011a. "Confronting Many-Many Problems: Attention and Agentive Control." Noûs 45 (1): 50–76.
- 2011b. "Attention as Selection for Action." In Attention: Philosophical and Psychological Essays, edited by Christopher Mole, Declan Smithies, and Wayne Wu, 97–116. New York: Oxford University Press.
- ——. 2011C. "What Is Conscious Attention?" Philosophy and Phenomenological Research 82 (1): 93–120.
- 2013a. "Mental Action and the Threat of Automaticity." In *Decomposing the Will*, edited by Tillman Vierkant, Julian Kiverstein, and Andy Clark., Oxford: Oxford University Press.
- -----. 2013b. "The Case for Zombie Agency." Mind 122 (485): 217-30.
- 2013c. "Book Review: The Conscious Brain: How Attention Engenders Experience, by Jesse Prinz. Mind. first published online November 28, 2013 doi:10.1093/mind/fzt101.
- Wurtz, R. H. 2008. "Neuronal Mechanisms of Visual Stability." Vision Research 48 (20): 2070–89.

308 BIBLIOGRAPHY

- Yantis, Steven. 1993. "Stimulus-driven Attentional Capture and Attentional Control Settings." Journal of Experimental Psychology. Human Perception and Performance 19 (3) (June): 676–81.
- Yantis, Steven, and John Jonides. 1984. "Abrupt Visual Onsets and Selective Attention: Evidence from Visual Search." *Journal of Experimental Psychology: Human Perception and Performance* 10 (5): 601–21.
- Yarbus, A. L. 1967. Eye Movements and Vision. New York: Plenum Press.

INDEX

Page numbers are grouped (e.g. 135–39) when the discussion of the term/ concept spans across consecutive pages; they are listed individually (e.g. 134, 135, 136) when the discussion of the term/concept is not continuous.

AIR (Attended Intermediate Representation) Theory 181, 183-86, 203, 204 Allport, A. 6, 29, 30, 41, 67-68, 74; on selection for action 76-80, 103 Anderson, B. 5 Anton-Erxleben, K. 140, 141, 142 Armstrong, K. 64, 66, 72 attention: bottom-up 26, 31, 34, 36, 91, 118–19, 142, 154, 165–70, 272; capture of 37-38, 43, 77, 91-93, 107, 256; control of 30-38; covert 64, 121, 133, 134, 152; feature 171, 210-11, 255; filter 12, 16-21, 41, 46, 57-58, 62, 71, 151-52, 162, 273; networks of 27-29; object 150,

174, 210-15, 236, 255; as orienting 27-28, 31, 36, 41, 43-44, 72, 92-93, 151-52; overt 23, 35, 36, 64, 151; see also multiple object tracking; partial 102-3, 252-53; as selection for action 77-79; spatial 23-29, 63-71, 113, 150, 173, 210-14; see also Premotor Theory; as structuring consciousness 129-31; top-down 26, 29-31, 34, 60, 119, 121, 166-67, 175, 183; to time 9, 237; unconscious 91, 111, 112-14, 126, 127, 130, 139, 140, 222-24, 234, 271; as vigilance 28, 93-95 automaticity 29-38, 82-83, 92-93, 107, 148-49, 238, 251-55

Baars, B. 181, 193, 204; see also Global Workspace Theory Baddeley, A. 105-6 Beck, J. 145 behavior space 79-90, 96-97, 106, 107, 249-51; see also Many-Many Problem Behrmann, M. 28 biased competition 48, 58-62, 70, 72, 273 blindsight 91, 112-14, 140, 141, 178, 222-24, 227-36, 245-46 Block, N. 140, 143; on access 178, 187, 190-203; on crowding 167-72, 175; against representationalism 134-40, 144, 145 Boolean Theory of Attention 226, 240-41 Boynton, G. 56, 61, 62, 196 brain areas: frontal eye field (FEF) 28, 53, 63-66; see also Premotor Theory; inferotemporal cortex (IT) 52-53, 59; lateral intraparietal sulcus (LIP) 53, 63; middle temporal area (MT/V5) 53, 113, 184; superior colliculus (SC) 63; visual area V1 52, 53, 112-13, 238; visual area V2 57, 113, 184; visual area V3 53, 113, 184; visual area V4 51, 53, 55, 57-58, 59, 63-66, 113, 184 Broadbent, D. 15-17, 21, 41, 58, 274, 279 Bronner, B. 125

Campbell, J. 111, 240, 246, 267; on attention and reference 225–27, 233–35; on selection and access 219–24 capacity limits 15-16, 19, 59, 77, 188, 198–99, 274, 292–93 Carandini, M. 62, 70 Carrasco, M. 26, 31, 41, 114-21, 124-26, 131-40, 141-42, 144 Carruthers, P. 267 Cavanagh, P. 168 Chalmers, D. 128, 131-33, 140, 143, 255, 264-65, 267, 269 change blindness 153-54, 159-60, 172, 193, 200 Cherry, C. 16 Chirimuuta, M. 175 Cognitive Unison Theory 99-103, 104, 108 Coltheart, M. 189 consciousness: access 177-78, 234; see also cognitive access; phenomenal 15, 111, 131, 140, 155, 178, 180-81, 186-87, 200, 205, 206, 208, 234; rational access 234 Corbetta, M. 42-43 Craver, C. 72 crowding 167-71, 175 Cumming, S. 108

Darden, L. 72 de Brigard, F. 4–5, 85 Dehaene, S. 181–83, 203, 204; *see also* Global Workspace Theory Desimone, R. 51, 57–59, 70, 151, 273 determinable/determinate distinction 125, 195–98, 207 Dicey Jennings, C. 104 dichotic listening 16, 38–39, 77, 105, 210, 237 Dickie, I. 230–32, 233, 236, 241 divisive normalization 48, 57–58, 61–63, 69–71, 74 Dretske, F. 137, 172 Driver, J. 16, 19, 25, 41, 212, 237 Duncan, J. 58–59, 211–14; see also biased competition

Egly, R. 212–14 empirical sufficient condition 39–40, 46, 62, 69–71, 72, 76, 83, 84, 96, 103, 112–14, 163, 167, 170, 214, 271 erotetic theory of attention 108–9 Evans, G. 236, 239, 240

Fallah, M. 63 Feature Integration Theory 21–23, 37, 41, 46, 48, 69–70, 148, 226, 273; see also A. Treisman Fechner, G. 115, 130, 141 FINST 218–19, 221, 231, 236, 239, 240; see also multiple object tracking Folk, C. 37–38, 43 Fuller, S. 120–21

gain modulation: contrast 56, 70; multiplicative 55, 69, 70 gabor patch 55, 117-19, 124, 125, 128, 129, 134-35, 141, 143, 144, 173, Ganson, T. 125 Gertler, B. 254, 262-69 Gibson, B. 37 gist 148, 164-67, 192, 199, 206 Global Workspace Theory 181-87, 203, 204, 207 Gobel, J. 119-20, 142 Goodale, M. 104–5, 161, 240 Groos, K. 4 Grush, R. 195 GY: patient 112-13, 141, 150-52, 222–23, 234, 246; see also blindsight Halligan, P. 74 Heeger, D. 61–63, 70; *see also* divisive normalization Helmholtz, H. 23 Huang, L. 226, 240–41

iconic memory 187–90, 194 inattentional agnosia 161–62, 172 inattentional amnesia 159–61, 164 inattentional blindness 146–72, 186, 187, 190, 202, 237, 274 information: mutual 14–16, 275–79 inhibition of return 27 introspection 114, 116, 191, 253–66

James, W.: on attention 3–4, 12–13, 64, 85, 91, 270–71; on attention and consciousness 110, 112, 115, 120, 130, 139–40, 141, 146, 233, 234–35, 273–74 Jonides, J. 37, 43 justification: doxastic 228, 230, 245–53, 268; propositional 228, 230, 232, 245–48, 267

Kahneman, D. 19–20, 225 Kelsey, E. 37 Kentridge, R. 112–13, 140, 150, 172 Klein, R. 64–65, 155 Koch, C. 15, 166–67, 172 Koralus, P. 104, 108–9; see also erotetic theory of attention Kouider, S. 191, 205–6

Lamme, V. 160, 189, 200–203, 207–8 Landman, R. 190, 200, 207 Lavie, N. 20–21, 41, 162; *see also* Load Theory
Lerman, H. 143 Load Theory 20-21, 38, 41, 162 Machamer, P. 72 Mack, A. 153-54, 160, 162, 172; on inattentional blindness and gist 164-67 Many-Many Problem 79-90, 98, 104, 107, 217, 249, 254 Marshall, J. 88 Maunsell, J. 55, 61, 62 mechanism 20, 45-46, 48, 49, 58, 72, 160, 170, 171; as answer to metaphysical question 67-71; of attention 27, 59-62, 63-66, 73, 74, 79, 185, 213, 218-22 Milner, D. 104-5, 161, 240 Mishkin, M. 105 Mole, C. 4, 41, 69-70, 74, 104, 105; on cognitive unison 99-103 Moore, T. 63-64, 66 Moran, J. 58, 151 Moray, N. 17 multiple object tracking 214-15, 218–19, 231, 236, 239–40 multiple realizability argument 67-69

Naccache, L. 181–83, 203, 204, 205 Nanay, B. 124–25, 136, 142 neglect 74–75, 154–55 Neisser, U. 189, 237 Neumann, O. 77–79, 103 Nickel, B. 133 Noë, A. 172

oculomotor readiness hypothesis 64; *see also* Premotor Theory Oliva, A. 164 overflow hypothesis 187, 189–94; alternatives to 192–200 Overgaard, M. 141

partial report 188-89, 193, 194, 198 Pashler, H. 55, 226, 240-41 Peacocke, C. 104, 247, 267 Peterson, S. 27-28 phenomenal salience 127-30, 141, 242 Phillips, I. 194, 203, 205, 207, 237 pop-out 22, 23, 26, 35, 37, 43, 272; versus capture of attention 91; and selection for action 92 Posner, M. 24-28, 39, 41, 49, 113, 150, 173-74; see also spatial cueing paradigm Premotor Theory 63-66, 72, 74, 152, 273 Prinz, J. 4-5, 141, 150-52, 172, 174, 181, 205, 207, 241; on AIR Theory 183-86, 203, 204 Pylyshyn, Z. 214-15; on FINST 218-8, 231, 236, 239-40

receptive field 51, 52, 55, 56, 59, 61, 63–65, 73, 138; olfactory 73; remapping of 54, 57–58, 60, 151–52; motor field 63, 65 relationalism 132, 134, 269 Remington, R. 37, 42 representational content 16, 111, 124, 129, 131–34, 135, 137–39, 143, 239, 259, 261 representationalism 131–34, 135–36, 139, 143 Reynolds, J. 61–62, 70 Rizzolatti, G 64–65; *see also* Premotor Theory

Rock, I. 153-54, 162, 165-66, 172, 237 Ruff, C. 72 Schacter, D. 85 Schenk, T. 66, 72, 105, 152 Schneider, W. 32-33 Scholl, B. 214-15, 219, 236 selection: early 14, 17, 20; early versus late 14-21, 41, 42, 162; late selection 14-21; for action 6, 9, 76-109, 110, 144, 163, 175, 185, 217, 218, 221, 233-34, 252, 271; for working memory 185 Shannon, C. 14-15, 57, 274; see also information theory Shapiro, L. 68 Shiffrin, R. 32-33 Shomstein, S. 28, 213, 236, 238 Shulman, G. 42-43 Siegel, S. 140. 245, 251-53, 267, 268 Silins, N. 245, 251-53, 267, 268 Simons, D. 160, 161, 172 Smithies, D. 227-30; see also consciousness, rational access Sosa, E. 255-56 spatial cueing paradigm 23-29, 39, 43, 49, 113, 150, 173, 212, 213 spatial filter 196-97 Speaks, J. 128, 140, 144 Sperling, G. 141, 187-89, 194, 198, 199, 201, 203, 205 spikes 46, 50, 51, 53, 59, 60 spotlight 3, 5, 22 Stazicker, J. 124, 125, 136, 140, 142, 163, 191, 196-97, 207 supervenience 67, 94, 107, 131-32, 134, 143 Sutherland, S. 4 synchrony 73-74

Taylor, H. 170, 188 Treisman, A. 17–18, 19–20, 42, 211; on attention as feature binding 41, 69, 148, 219, 273; on Feature Integration Theory 21–23, 48, 219, 273 Tsal, Y. 20, 41 Tse, P. 121–22, 142 Tsotsos, J. 41, 48

Ungerleider, L. 60, 105

visual acuity 142, 151, 168 visual search 20, 21–23, 35, 37, 38, 41, 43, 83, 85, 105, 108, 210–11, 272

Ward, L. 26, 36, 41, 72 Watzl, S. 125, 144, 268; on attention as structuring consciousness 129-30, 141, 143 Weaver, W. 42, 278 Wolfe, J. 4, 41, 56, 160, 161 working memory 15, 27, 74, 84, 105-6, 151-52, 188, 189, 202, 204-5, 207, 242; AIR Theory and 183-86; capacity of 191-92, 195, 198-200; cognitive access and 176-81; Global Workspace as 181-83; inattentional amnesia and 161 Wright, R. 26, 36, 41, 72 Wu, W. 33, 79, 125, 132, 140, 174, 204, 268; on Many-Many Problem 82; on attention as selection for action 104, 105; on phenomenal salience 127–29

Yantis, S. 37, 43, 213, 238 Yarbus, A. 35–36, 43, 83 Yeshurun, Y. 119

