The Reflexive Nature of Consciousness

Greg Janzen

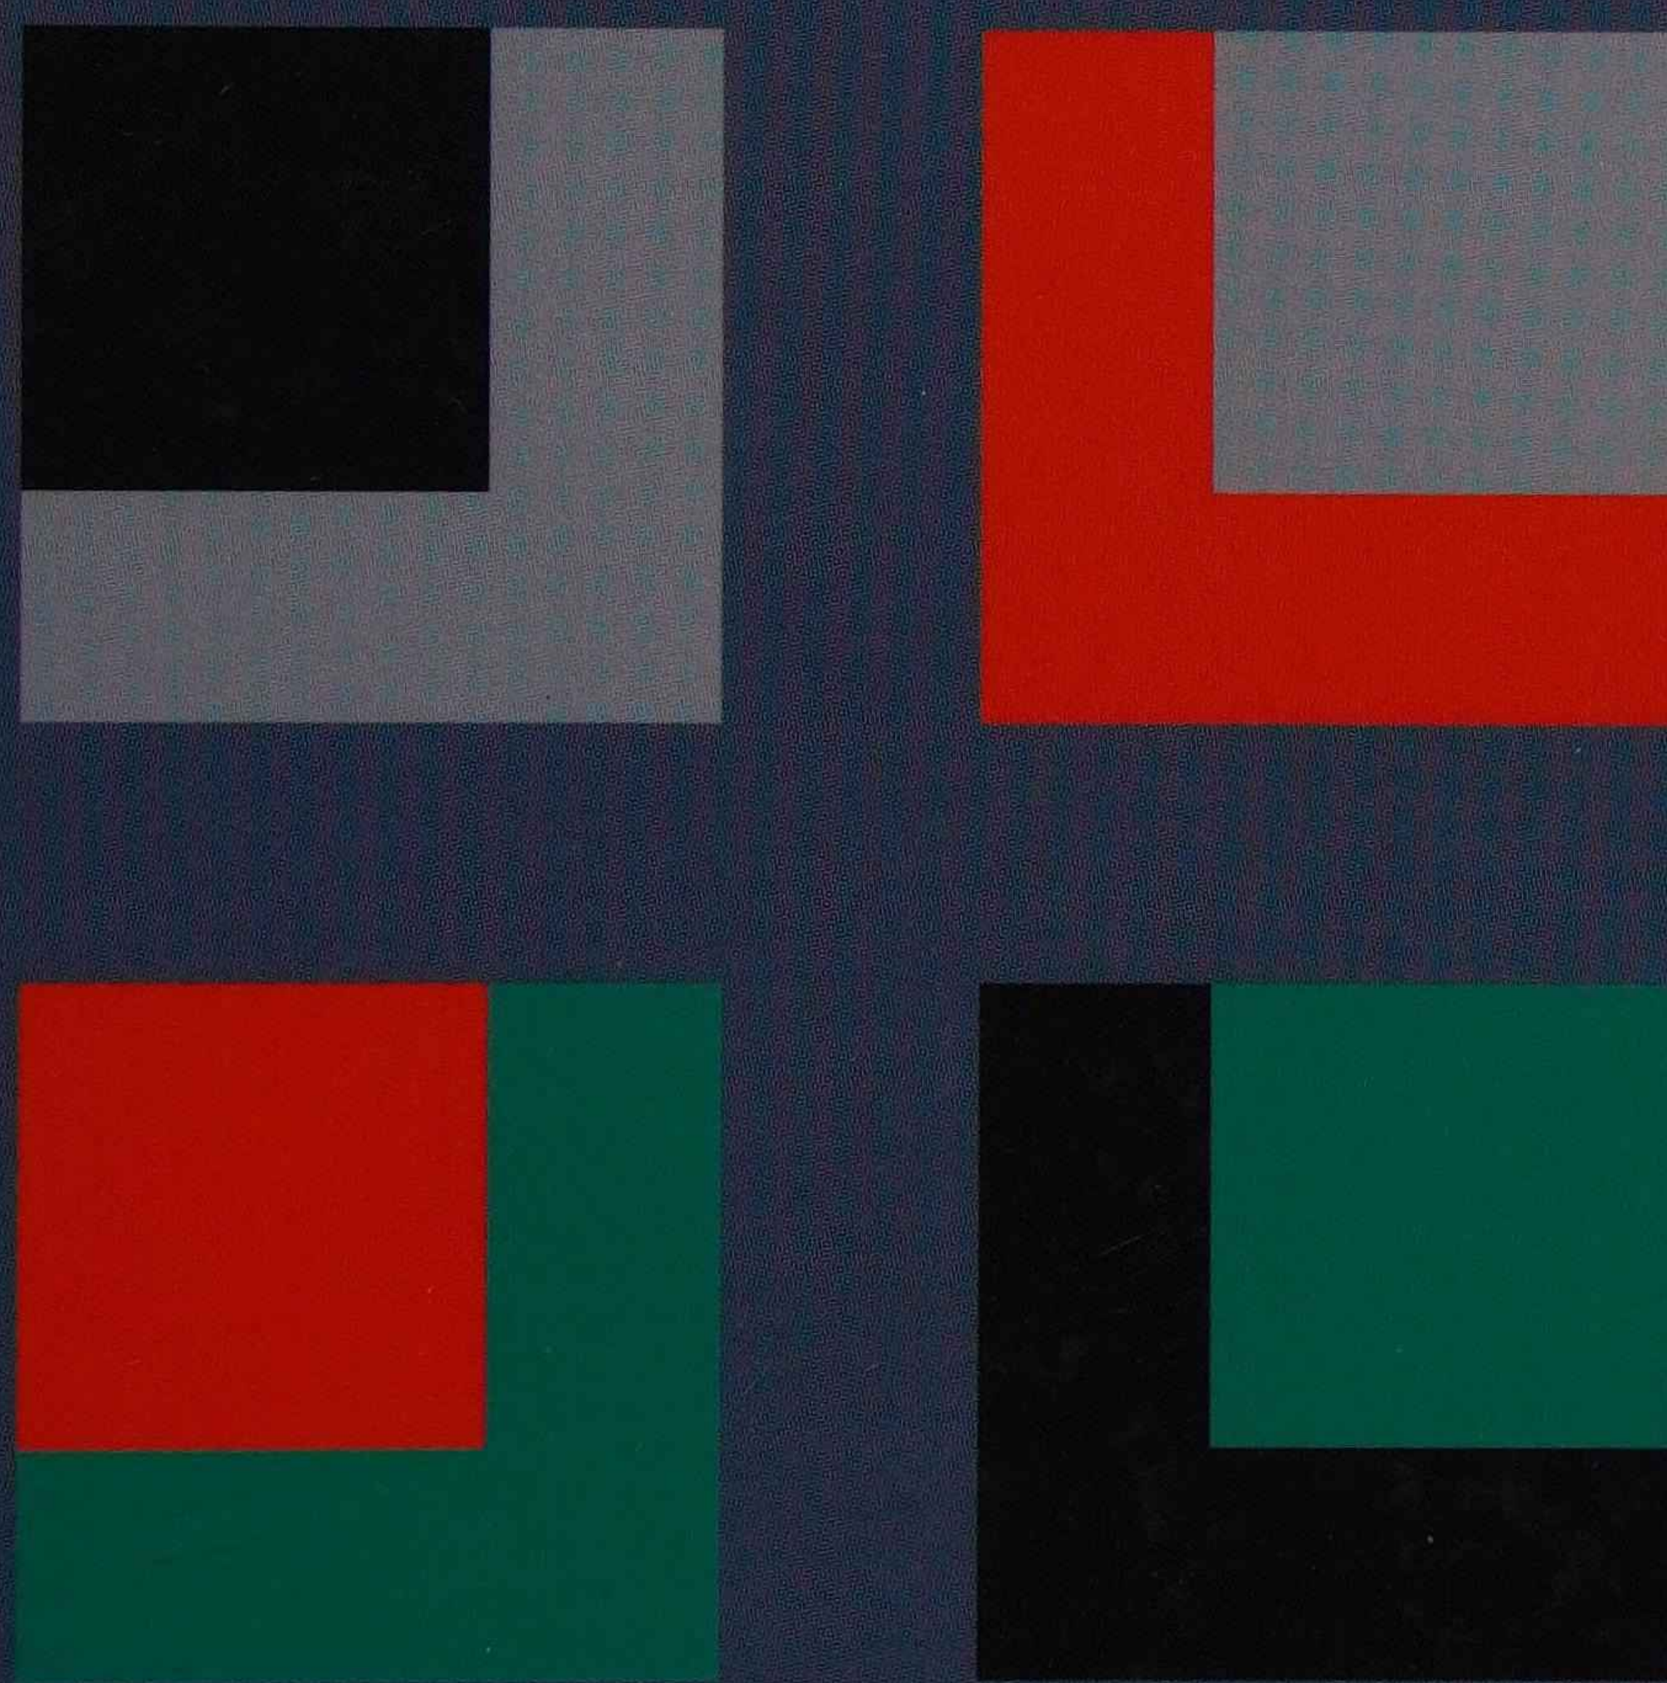

Advances in Consciousness Research
72

John Benjamins Publishing Company

LIVERPOOL
JOHN MOORES UNIVERSITY
AVRIL ROBARTS LRC
TEL 0151 231 4022

WITHDRAWN

126

JAN

liverpool JMU
Learning and Information Services

Accession no D674821 ω	
Supplier DA	Invoice Date
Class no	
Site V	Fund code PSY6

Books are to be returned on or before
the last date below.

2·12·2010

LIVERPOOL
JOHN MOORES UNIVERSITY
AVRIL ROBARTS LRC
TEL 0151 231 4022

LIBREX

LIVERPOOL JMU LIBRARY
3 1111 01268 8774

The Reflexive Nature of Consciousness

Advances in Consciousness Research (AiCR)

Provides a forum for scholars from different scientific disciplines and fields of knowledge who study consciousness in its multifaceted aspects. Thus the Series includes (but is not limited to) the various areas of cognitive science, including cognitive psychology, brain science, philosophy and linguistics. The orientation of the series is toward developing new interdisciplinary and integrative approaches for the investigation, description and theory of consciousness, as well as the practical consequences of this research for the individual in society.
From 1999 the Series consists of two subseries that cover the most important types of contributions to consciousness studies:
Series A: Theory and Method. Contributions to the development of theory and method in the study of consciousness; Series B: Research in Progress. Experimental, descriptive and clinical research in consciousness.
This book is a contribution to Series A.

Editor

Maxim I. Stamenov
Bulgarian Academy of Sciences

Editorial Board

David J. Chalmers
Australian National University

Gordon G. Globus
University of California at Irvine

Christof Koch
California Institute of Technology

Stephen M. Kosslyn
Harvard University

Stephen L. Macknik
Barrow Neurological Institute, Phoenix, AZ, USA

George Mandler
University of California at San Diego

Susana Martinez-Conde
Barrow Neurological Institute, Phoenix, AZ, USA

John R. Searle
University of California at Berkeley

Petra Stoerig
Universität Düsseldorf

Volume 72

The Reflexive Nature of Consciousness
by Greg Janzen

The Reflexive Nature of Consciousness

Greg Janzen
University of Calgary

John Benjamins Publishing Company
Amsterdam / Philadelphia

The paper used in this publication meets the minimum requirements of American National Standard for Information Sciences – Permanence of Paper for Printed Library Materials, ANSI Z39.48-1984.

Library of Congress Cataloging-in-Publication Data

Janzen, Greg.
The reflexive nature of consciousness / by Greg Janzen.
p. cm. (Advances in Consciousness Research, ISSN 1381-589X ; v. 72)
Includes bibliographical references and index.
1. Consciousness. 2. Phenomenology. 3. Self-consciousness (Awareness) 4. Self. I. Title.
B808.9.J36 2008
126--dc22 2007045723
ISBN 978 90 272 5208 1 (Hb; alk. paper)

© 2008 – John Benjamins B.V.
No part of this book may be reproduced in any form, by print, photoprint, microfilm, or any other means, without written permission from the publisher.

John Benjamins Publishing Co. · P.O. Box 36224 · 1020 ME Amsterdam · The Netherlands
John Benjamins North America · P.O. Box 27519 · Philadelphia PA 19118-0519 · USA

Table of contents

Acknowledgements VII

1. Introduction 1
   1.1 Conscious states, reflexivity, and phenomenal character 1
   1.2 Some remarks on methodology 4
      1.2.1 The primacy of phenomenology 5

2. Some semantics of "consciousness" 11
   2.1 Preamble 11
   2.2 Creature consciousness: transitive and intransitive 12
   2.3 State consciousness 16
   2.4 Unconscious mental states 21
   2.5 Self-consciousness 22
   2.6 Phenomenal consciousness 23
      2.6.1 Hacker and Lycan's criticisms of the "what-it-is-like" locution 25

3. A formula for state consciousness 33
   3.1 Nagel's what-it-is-like formula 33
   3.2 Putative counterexamples 34
      3.2.1 Access consciousness and phenomenal consciousness 37
      3.2.2 The daydreaming, long-distance truck driver 46
   3.3 Non-conscious phenomenality? 53
   3.4 Summary 56

4. Consciousness and self-awareness 59
   4.1 Preamble 59
   4.2 A gloss on intentionality 59
   4.3 The Transitivity Principle 69
   4.4 Two positive arguments for the Transitivity Principle 74
      4.4.1 The Symmetry Argument 75
      4.4.2 Objections and replies 78
      4.4.3 The argument from spontaneous reportability 79

5. **Higher-Orderism** 85
5.1 Preamble 85
5.2 The Higher-Order theory of consciousness 86
5.2.1 Confabulated conscious states 96
5.2.2 A phenomenological objection 99

6. **A "one-state" alternative** 103
6.1 Preamble 103
6.2 The Brentanian model 104
6.2.1 The regress argument 110
6.2.2 A twist 115
6.2.3 *De Se* content 118
6.2.4 The "self" 129
6.3 Objections and replies 131

7. **Representationalism** 137
7.1 Preamble 137
7.2 The representational theory of phenomenal character 137
7.2.1 The Transparency Assumption 141
7.2.2 The Property Assumption 147

8. **The nature of phenomenal character** 155
8.1 Preamble 155
8.2 Phenomenal character as implicit self-awareness 156
8.2.1 Differences in phenomenal character 163
8.2.2 Pains and other bodily sensations 166
8.3 Some brief remarks on privacy 169
8.4 Conclusion 172

Bibliography 173
Index of subjects 183
Index of names 185

Acknowledgements

A number of people have read and commented on (ancestors of) sections of this work. I owe special thanks in this regard to John Baker. Thanks also to Phil Dwyer and Miri Albahari for many helpful conversations. My thanks also to an anonymous reviewer for John Benjamins Publishing, whose criticisms and suggestions occasioned a number of substantial improvements.

Some of this material has been previously published. A large portion of Chapter Three and part of Chapter Four originally appeared in "Self-Consciousness and Phenomenal Character," *Dialogue* (2005) **44**: 707–33. Some of Chapter Six and some of Chapter Eight originally appeared in "Phenomenal Character as Implicit Self-Awareness." *Journal of Consciousness Studies* (2006) **13**: 44–73. I thank those journals for reprint permission. The bulk of Chapter Seven is a revised and expanded version of "The Representational Theory of Phenomenal Character: A Phenomenological Critique," *Phenomenology and the Cognitive Sciences* (2006) **5**: 321–39. It is reprinted here with kind permission from Springer Business and Media.

CHAPTER 1

Introduction

1.1 Conscious states, reflexivity, and phenomenal character

A central problem in contemporary discussions of consciousness is the problem of determining what a mental state's being conscious consists in. This book presents a distinctive approach to this problem, arguing that a mental state is conscious if and only if it has a certain reflexive character, i.e., if and only if it has a structure that includes a built-in awareness (or consciousness) of its own occurrence.[1] Since this thesis finds one of its clearest expressions in the work of Brentano, it is his treatment of the thesis on which I initially focus, though I subsequently bring in Sartre where he is required to improve on Brentano, i.e., where he addresses himself to an important point not considered by Brentano.[2]

Another, more specific, aim of this investigation is to shed some light on the phenomenal, or "what-it-is-like," dimension of conscious experience. Despite the burgeoning literature on consciousness, relatively little enlightening has been said about the nature of phenomenal character, an integral – indeed, as I will argue, *defining* – feature of consciousness. This book attempts to remedy this lack, arguing that the relationship between phenomenal character and self-awareness has so far gone under-recognized in the literature on consciousness.[3] I attempt to show, in particular, that the phenomenal character of conscious

---

1. Throughout this book I will regard "consciousness" and "awareness" (and their cognates) as equivalent. Strictly speaking, these terms are not synonymous – e.g., one can be unconsciously aware of *x*, but one cannot be unconsciously conscious of *x* – but nothing will be lost if I use them interchangeably. For further discussion of the differences between consciousness and awareness, see Bennett and Hacker 2003, p. 256; Gennaro 1996, pp. 5–7; and A. R. White 1964, ch. 4.

2. To be clear, however, my concern is with Brentano's reflexivity thesis, not with Brentano. This book is not an exercise in historical commentary.

3. A comment on my use of "self-awareness" should be made at the outset. Some think that awareness of one's own mental states is not *self*-awareness unless the subject, in being conscious of *x*, is aware of *herself* as being conscious of *x*. Properly speaking, this is true, but I will occasionally use "self-awareness" rather loosely; that is, I will sometimes use the term to denote a subject's being aware just of her mental state. The context should make clear what is being intended.

states is constituted by a type of self-awareness, i.e., by a low-level or "implicit" self-awareness that is built into every conscious state.[4] This thesis, if accepted, bodes well for consciousness studies, since it is widely supposed, among both philosophers and scientists, that what makes the problem of consciousness so intractable is the phenomenal character of experience (e.g., Nagel 1974; McGinn 1991a; Chalmers 1996; Edelman 1989). But if phenomenal character is constituted by implicit self-awareness, then the problem of phenomenal character *just is* the problem of implicit self-awareness. Considerations of simplicity and unification, then, seem to favour the sort of analysis of phenomenal character I am proposing.

Contemporary philosophical discussions of consciousness tend to take off from Nagel's seminal and much-anthologized paper "What is it Like to Be a Bat?" (1974), in which Nagel argues that all conscious states have a distinctive phenomenal character, or *something it is like* to have them (cf. Farrell 1950; Sprigge 1982). In general, however, the ensuing debate has revolved around the question of whether phenomenal character poses a threat to physicalism (or materialism), a well-worn debate to which I have little to contribute.[5] My aim in these pages is not to advance the debate over whether consciousness can be explained in purely physical terms, but rather to provide an answer to the preliminary question of what it is for mental state to be conscious in the first place, and, in so doing, provide an account of the nature of phenomenal character. I will attempt to show that a philosophically profitable approach to these matters is, perhaps surprisingly, to view conscious states as reflexive. This approach to the analysis of state consciousness has antecedents in the writings of, among others, Aristotle, Sartre, Gurwitsch, and, as I remarked, Brentano, but is often neglected in current discussions. Moreover, even those philosophers sympathetic to the reflexive approach have not developed it in any kind

---

4. To say that the relationship between phenomenal character and self-awareness has gone under-recognized is not to say that it has not been recognized at all. Flanagan (1992), Zahavi (1999, 2003, 2004), and Zahavi and Parnas (1998), for example, have characterized phenomenal consciousness in terms of self-awareness, and Kriegel (2003a) has argued that a mental state is conscious just in case it involves implicit self-awareness. As far as I know, however, nobody has defended the claim that the phenomenal character of conscious states is *constituted by* self-awareness.

5. Crane and Mellor (1990) have made an awfully compelling case for the claim that the ceaseless disputes over mental states are physical or non-physical is a hopeless red-herring. "Physicalism," they argue, "is the wrong answer to an essentially trivial question.... It cannot begin to help philosophers of mind answer the serious questions about the mind.... [for example] what enables some parts of the world (us) to [be conscious] of other parts, including other people (and of course ourselves)" (1990, p. 206; cf. McGinn 2003, p. 162n17).

of substantive detail.[6] My aim is to inform these debates by providing a detailed account of the reflexive approach and to assess its viability against non-reflexive approaches.

In the remainder of this chapter I provide some remarks on methodology; specifically, I spell out an important methodological principle that has guided the philosophical view, or set of views, propounded in this book, namely, the principle of the primacy of phenomenology.

In Chapters Two and Three I examine the concept of state consciousness, arguing that, contrary to accepted belief, the concept admits of a relatively straightforward analysis in terms of necessary and sufficient conditions. More fully, in Chapter Two I examine a number of distinctions frequently invoked in discussions of consciousness, including transitive creature consciousness, intransitive creature consciousness, state consciousness, self-consciousness, and phenomenal consciousness. This precipitates a discussion of phenomenal character, in which I defend the claim that phenomenal character is best understood in terms of the "what-it-is-like" locution. Then, in Chapter Three, I advance considerations suggesting that a subject is in a conscious state just in case that state has phenomenal character, i.e., just in case there is something it is like for the subject to be in that state.

In Chapter Four, I commence my analysis of the relationship between state consciousness and self-awareness, arguing, first, that all conscious states have intentionality (i.e., are of or "about" something) and, second, that all state consciousness involves self-awareness, in the sense, roughly, that a conscious mental state is always a state of which the subject is, in some appropriate way, aware. My arguments for this latter thesis – which, following Rosenthal (e.g., 2000b, 2005), I call the Transitivity Principle – exploit the assumption, defended in Chapter Three, that conscious experiences are necessarily like something to the subject of the experience.

In Chapter Five I examine a currently popular theory of consciousness that takes conscious states to necessarily involve self-awareness, viz., the higher-order theory, according to which state consciousness is explicable in terms of a relation obtaining between a lower-order mental state and an actual or potential higher-order thought or experience. The upshot of my analysis of the higher-order theory is that it is both logically and phenomenologically inadequate.

---

6. Though there is what might loosely be called a reflexive approach to consciousness that *has* received rigorous theoretical treatment recently, namely, the higher-order approach (e.g., Rosenthal 1986, 1990; Gennaro 1996; Lycan 1996). This approach might loosely be called reflexive because it takes conscious states to necessarily involve self-awareness. However, as I make clear in Chapters Five and Six, the higher-order approach to consciousness is importantly different from the sort of reflexive approach Brentano espoused.

Further pursuing the idea that state consciousness is intimately bound up with self-awareness, in Chapter Six I explicate the Brentanian model of consciousness, which, unlike higher-order models, takes self-awareness to be an intrinsic feature of conscious states, in the sense that a conscious state has a structure that includes an awareness of itself. However, the Brentanian model requires revision, since, just like higher-order models, it erroneously construes self-awareness in terms of a "subject-object" relation, i.e., in terms of conscious states being objects to which subjects stand in a certain kind of relation. Later in Chapter Six, therefore, I propose an alternative to higher-orderism that combines Brentano's central insight – i.e., that there is a single token state involved in one's seeing *x*, hearing *x*, etc. – with Sartre's claim that even though our conscious states themselves are always "given" or "manifested" in experience, they do not become *objects* of consciousness except during acts of reflection. In the final section of this chapter, I defend this Brentanian/Sartrean model of experiential self-givenness against some possible objections.

In Chapter Seven, I return to the issue of phenomenal character, providing an extended discussion of the representational theory of phenomenal character. According to the representational theory, the phenomenal character of an experience is entirely determined by, and is in fact identical with, the experience's representational content. I show that the representational theory cannot withstand careful scrutiny: when it is laid bare, it is found to rest upon a pair of dubious assumptions.

Finally, in light of my discussion of the representational theory, I defend, in Chapter Eight, the claim that phenomenal character is constituted by an implicit self-awareness that is built into every conscious state.

1.2 Some remarks on methodology

As I mentioned, in writing this book I have closely adhered to an important methodological principle: the principle of the primacy of phenomenology. Since this principle guides the philosophical view advanced in this book, it seems appropriate that I should provide a brief statement of it. It bears emphasizing that my statement of the principle is, for the most part, just that, a statement of it, and not any kind of rigorous defence of it. The plausibility of the principle (or some variation thereof) has, I believe, been ably demonstrated by others,[7] and so it has, for me, the

7. I am referring, of course, to the classical phenomenologists, such as Brentano (1874), Husserl (1960), Merleau-Ponty (1962), and Sartre (1956). The principle has also been defended by – or would (I expect) find a natural ally in the work of – a number of contemporary philosophers, including Nagel (1974), Woodruff Smith (e.g., 1989), Dwyer (1990), Siewert (1998), Stubenberg (1998), Dainton (2000), Thomasson (2000), Zahavi (e.g., 1999, 2005), and Carman (2005).

status of a datum from which theorizing about consciousness should begin. Thus, I will simply state the principle and proceed accordingly.

1.2.1 The primacy of phenomenology

Along with virtually every major figure in phenomenology, Nagel maintained that a non-negligible feature of consciousness is its "first-personal" or "subjective" dimension – i.e., the way it is experienced from the first-person perspective. He stressed that if we neglect the first-personal perspective in favour of an exclusively "third-personal" or "objective" point of view, we will inevitably abandon the essential feature of consciousness. In accordance with this assumption, and the assumption that the chief aim of phenomenology is to provide a description of the first-personal dimension, the principle of the primacy of phenomenology states, basically, that any account of consciousness is apt to founder unless it is phenomenologically sensitive. The term "phenomenology" is frequently used nowadays to denote conscious states that have phenomenal character, but phenomenology, at least if understood in its more "technical" sense, i.e., in terms of the philosophical movement that began as Brentano's descriptive psychology and developed in the work of Husserl, is the *act* (more or less) of reflecting on and describing one's conscious states. It is important, however, to guard against a widespread misconception, viz., that phenomenology is a kind of pre-theoretical, "folk psychological" method of philosophizing about consciousness. Zahavi and Parnas have noted that philosophers working within the analytic tradition tend to "tacitly identify phenomenology and the commonsense considerations of 'folk psychology'" (1998, p. 688; cf. Gallagher and Varela 2003), but such an identification vastly oversimplifies the aims and scope of the phenomenological method. While phenomenology is, at least in part, a matter of what Cowley (1968) calls "reflexive analysis," i.e., a matter of reflecting on and describing our naïve perceptual and non-perceptual experience,[8] it is also a matter of doing so with an eye toward recognizing the general experiential structure of perceptual and non-perceptual consciousness. As Merleau-Ponty has put it: "Phenomenology is the study of essences; and according to it, all problems

---

8. This is not quite right, at least for Husserl. For Husserl, the proper method of phenomenology involves initiating an *epoché* whereby one "brackets" or places in abeyance one's naïve perceptual experience (see Husserl 1960). According to Husserl, the *epoché* discourages prejudices that might otherwise hinder the discovery of the "universal structures" underlying conscious experience. Merleau-Ponty (1962), on the other hand, baulked at the notion of an *epoché*, arguing that our naïve perceptual experience is precisely what phenomenology does describe. For Merleau-Ponty, to use Cowley's words, "there is no sort of sense-experience other than that of the vulgar consciousness" (Cowley 1968, p. 19).

amount to finding definitions of essences: the essence of perception, or the essence of consciousness, for example" (1962, p. vii). More recently, Zahavi, echoing Merleau-Ponty, has put the point this way: "Phenomenology is interested in the very dimension of givenness or appearance and seeks to explore its essential structures and conditions of possibility" (2007, p. 31). Phenomenology, then, is not concerned with folk-psychological reports of subjects' occurrent experiences; on the contrary, it aims to disclose the general experiential structure of conscious experience. And it is reasonable to suppose that a careful phenomenological analysis can provide a more accurate model of the general experiential structure of conscious experience than the models (read: computational/functional, neurophysiological, etc.) that are currently so popular in philosophy and cognitive science.

Phenomenology can provide a more accurate model for the simple reason that, in an important sense, *only* phenomenology can provide a model. A fallacy that has become distressingly prevalent in recent attempts to understand consciousness involves supposing that the way a conscious experience is in itself might transcend the way it is for the experiencing subject, i.e., the way phenomenology reveals it to be.[9] For example, in response to Edelman's claim that "[o]ne of the most striking features of consciousness is its continuity" (1989, p. 119), Dennett makes the following claim:

> This is utterly wrong. One of the most striking features about consciousness is its discontinuity – as revealed in the blind spot, and saccadic gaps, to take the simplest examples. The *dis*continuity of consciousness is striking because of the *apparent* continuity of consciousness.
>
> (1991, p. 356)

This claim has dramatic implications: if Dennett is right, then perceptual consciousness, as it really is, is irregular and sporadic. The continuity we experience in naïve perception, one is given to understand, is merely "apparent," an illusion. The upshot: conscious perceptual experience is one way in itself and another way for the experiencing subject, and the way it *really* is, i.e., the way it is in itself, is often quite distinct from how it is for the experiencing subject.

---

9. For interesting and stimulating recent discussions of this fallacy see Dwyer (1990, passim) and Rowlands (2001, ch 8, §7ff.). Both Dwyer and Rowlands characterize the fallacy as a category mistake. According to Dwyer, subjectivity encompasses "categorially autonomous phenomena" (1990, p. 33), and so it is a fallacy to suppose that our actual experience "may be some way in itself transcending the way it is for a subject" (1990, p. 32); and Rowlands claims that since personal and sub-personal levels of content cannot be reckoned to be in the same category, it is a fallacy to suppose that there is a "distinction between (i) the way an experience seems to its subject and (ii) the way an experience really is" (2001, p. 189).

But this is profoundly misconceived. To suggest that conscious experience, perceptual or otherwise, might transcend the way it is for the experiencing subject is to conflate two categorially distinct levels of content: the *personal* and the *sub-personal*.[10] At the sub-personal level, we have various neurological (cognitive, etc.) processes that are, presumably, necessary for experience. At the personal level, by contrast, we have the "field of experience": this is the level of content that phenomenology describes, the level of content that we experience from the first-person perspective.[11] And no sub-personal consideration can displace or supersede this level of content.[12] Thus, it makes no sense to say that perceptual consciousness, despite its apparent continuity, is *really* marked by discontinuity, just as it makes no sense to say that pain, despite its apparent painfulness, is not *really* painful. There is no such appearance/reality distinction to be made with respect to conscious experience. As Searle has succinctly put it: "where the existence of conscious states is concerned, you can't make the distinction between appearance and reality, *because the existence of the appearance is the reality in question*" (1997, pp. 121–2).[13] This is not to deny that the neurological, cognitive, etc. structures that are causally related to conscious experience (e.g., the neurological structures that are indispensable for seeing blue) merit investigation; and this is not to deny that we should endeavour to provide a phenomenological framework in which the phenomenological properties recognized are made continuous with the properties admitted by neuroscience, cognitive science, etc. However, it is to deny that sub-personal considerations somehow take precedence in the study of consciousness, i.e., that sub-personal considerations can give us insight into the "real" nature of

---

10. I am not entirely confident that the charge of conflating these different levels of content can be levelled against Dennett, since, in his 1969, he goes to significant lengths to distinguish between personal and subpersonal levels of explanation, where the former focuses on a subject's conscious interaction with his or her environment and the latter focuses on the physical-cum-cognitive enabling conditions of such interaction. But a case can be made, I think, that Dennett now believes, even if he once didn't, that subpersonal explanations of conscious phenomena ultimately eclipse any personal ones. For more on this, see below in the text.

11. The wording here is perhaps somewhat infelicitous. Strictly speaking, we do not experience a level of content; rather, we experience objects, events, properties, states of affairs, etc. The level of content described in phenomenology just is our experience of these worldly items.

12. Evident here will be the tangential inspiration of Putnam, who, in his 1987 (among other places), contends that we must abandon the view that what is *real* about the mind is somehow "behind" or "more fundamental than" ordinary appearances.

13. See also Dwyer (1990, pp. 32–3), McGinn (1995, p. 247), Rowlands (2001, pp. 189–94), Nagel (1974, p. 448), and Stubenberg (1998, p. 38).

consciousness, whereas the experiential considerations phenomenology describes cannot. Phenomenology, in other words, antecedes the empirical sciences of the mind: the data these disciplines provide must be made to cohere with the phenomenological data, not vice versa. There is a sense in which, tactically speaking, we must, as Flanagan says, "examine the fit between [the neuroscientists', cognitive scientists', etc.] stories and the phenomenological…stories" (1992, p. 11). But the appropriate "direction of fit" is, surely, phenomenology → neuroscience, cognitive science, etc. "Our understanding [of consciousness]," to quote Woodruff Smith, "will *begin* with the structure of consciousness analyzed in phenomenology" (2004a, p. 37, my emphasis). One upshot here is that the question of whether various conscious phenomena (perceptions, pains, emotions, etc.) can be preserved in our best neuroscientific, cognitive scientific, etc. account of consciousness does not arise, since, given their status as phenomenological data (and not theoretical constructs), these phenomena *have* to be preserved.

Now Dennett would no doubt demur at this claim, since he holds that there is no such thing as *actual* phenomenology (1991, p. 365), there only *seems* to be: "*There seems to be phenomenology*. That's a fact the heterophenomenologist enthusiastically concedes. But it does not follow from this undeniable, universally attested fact that *there really is* phenomenology" (1991, p. 366). Dennett defines phenomenology "as the generic term for the various items in conscious experience that have to be explained" (1991, p. 45), and so in saying that there is no actual phenomenology, Dennett appears to be suggesting, in an eliminativist manner, that there are no actual items in conscious experience that have to be explained. But Dennett is not, strictly speaking, an eliminativist. He holds that we are conscious, but that consciousness lacks various first-personal properties (sensations, feels, etc.) it is typically thought to have. Our experiences, on Dennett's view, are constituted solely by our *judgements* or *beliefs* about them: "There is no such phenomenon as really seeming – over and above the phenomenon of judging in one way or another that something is the case" (Dennett 1991, p. 364).[14] But the claim that there are no genuine experiential properties, only "seeming" experiential properties, is altogether opaque. Perhaps the idea is that just as an extra-mental

14. This is, arguably, a *kind* of eliminativism, but it is nowise akin to the full-fledged eliminative materialism that forthrightly denies the existence of consciousness altogether. This kind of eliminativism has some acute thinkers among its adherents (e.g., Wilkes 1984; Rey 1988), but has always struck me as a manifestly untenable philosophical position. I am not interested in providing any sort of refutation of it here, however, so I will simply defer to others who have. See, for example, McGinn 1989, ch. 2; Boghossian 1990; Searle 1992, pp. 46–9; Goldman 1993; and, in particular, Hacker 2001.

object can seem a certain way to an observer, so a mental state can seem a certain way to its bearer; only, there is nothing over and above the subject's judgement or belief regarding how his mental state seems to him. But this comparison is hardly illuminating, for although it makes sense to say that it seems to one as though yonder building is a barn even though yonder building is not really a barn, it makes no sense to say that it seems to one as though one is in pain even though one is not really in pain. It's part of the very concept of pain (and of our sensation concepts generally) that if it seems to one as though one is in pain, then one is.[15]

To make this point more perspicuous, it will be useful to consider Hill's (2004) recent attempt to show that there is an appearance/reality distinction to be made with respect to pain. According to Hill, it is possible to distinguish between its seeming to one as though one is pain and its being the case that one is in pain, since there is evidence that can override a subject's sincere avowal that he is in pain. For example, if a doctor tells a patient that he is perfectly healthy, the patient might come to believe that his chronic pain is psychosomatic. But this misses the point. To say that there is no appearance/reality distinction to be made with respect to pain is not to say that psychosomatic pain is impossible. And it is not to deny that a subject's pain might subside after he learns from a trusted authority that he is healthy. The point is that if it *actually* seems to one as though one is in pain, then one is in pain. It follows that some psychosomatic pain is genuine pain (since there is nothing incoherent about supposing that it could actually seem to one as though one is in pain even though one's pain is psychosomatic), but this consequence should not be troubling. It is possible for a subject to continue to experience (agonizing) pain despite coming to learn, and sincerely believe, that his pain is psychosomatic. And the subject's psychosomatic pain, in such cases, is no less real, i.e., no less that which disposes the subject to seek relief (perhaps through psychotherapy), than pain that is caused by injury or disease.

The principle of the primacy of phenomenology is extremely important, for it places constraints on what counts as a plausible account of consciousness and, correlatively, phenomenal character. Since any adequate account of consciousness must be phenomenologically sensitive, a necessary condition for an acceptable account of consciousness is that it be consistent with the principle of the primacy of phenomenology. One of my aims in this book is to develop an account that satisfies this condition.

---

15. Also see Shoemaker (1996b, pp. 226–9) and Hossack (2002, p. 170).

CHAPTER 2

Some semantics of "consciousness"

2.1 Preamble

Those who write about consciousness nowadays almost invariably begin by saying that the concept of consciousness is not a univocal concept admitting of necessary and sufficient conditions, but rather a slippery and mutable concept that signifies a heterogeneous class of psychological phenomena.[1] Consider the following remarks:

> While one should be wary of claims that important terms are systematically ambiguous, "consciousness" really is systematically ambiguous.
>
> (Nelkin 1993, pp. 419–20)

> As with most words, it is not possible to give a definition of "consciousness" in terms of necessary and sufficient conditions, nor is it possible to define it in the Aristotelian fashion by way of genus and differentia.
>
> (Searle 1992, p. 83)

> [Consciousness] can be frustratingly diaphanous: in talking about conscious experience, it is notoriously difficult to pin down the subject matter.
>
> (Chalmers 1996, p. 3)

> Nobody, I assume, would wish to argue that 'conscious', or its nominalisation 'consciousness', are clear notions.
>
> (Wilkes 1984, p. 224)

> The term 'consciousness' is multiply ambiguous.
>
> (Kriegel 2004, p. 182)

> [T]he philosophical use of "conscious" is by now well and truly up for grabs.
>
> (Lycan 1996, p. 25)

These writers are often drawn to the conclusion that "consciousness" is ambiguous because of the complex grammatical structure underlying the ordinary use of "consciousness" and its cognates, a complex grammatical structure that, upon

---

1. A few philosophers have tried to undermine the consensus that "consciousness" is ambiguous. The dissenting voices include McBride 1999 and Antony 2001. McBride has advanced a "unitary thesis," according to which the "verbal distinctions involving use of the word 'conscious' are unified in their reliance on a single ontological unit, that of conscious experience" (1999, p. 181). And Antony has argued that, within the consciousness literature, "consciousness" is univocal.

closer inspection, allegedly reveals a correspondingly complex psychological structure. The result, typically, is "a judicious division of labor" (Rosenthal 1990, p. 729), motivated by our linguistic practices, in which consciousness is partitioned into various types and subtypes.

In this chapter and the next, I review this judicious division of labour, advancing the claim that, despite the existence in the vernacular of different uses of the word "conscious," the verbal distinctions involving use of the term, at least as the term applies to states (or experiences),[2] unvaryingly depend on a specific property, that of *phenomenality* or *phenomenal character*. That is to say, "conscious state" and its cognates denote, unambiguously, states that have qualitative or phenomenal content, and the marker of the possession of this content is that there is *something it is like* for the subject of the state to be in it. Thus, phenomenal character is an essential property of conscious states, which means that when theorists speak of conscious states that lack phenomenal character they betray their confusion, speaking not of *conscious* states, but of something else. To be sure, the terms "essential" and "property" are metaphysically worrisome, and doubly so when taken together. But I share the position that there are bona fide properties or characteristics that some things, at least, possess, and that these properties are, for lack of a better word, essential to that thing, defining what it is. Phenomenal character, accordingly, is an essential property of conscious states; so no phenomenal character, no conscious state. The end result is that the apparently diffuse and multifarious concept of a conscious state actually has a relatively clear-cut underlying logical structure.

The following distinctions, or variants thereof, appear frequently in the literature. I will assume these distinctions, modulo, of course, the comments to follow in this chapter.

2.2 Creature consciousness: transitive and intransitive

"Conscious" is often used as an adjective of creatures, as in: "Smith is conscious after being in a coma for two days," or "She wasn't sure whether Smith was conscious." In these cases, "conscious" functions as an adjective that modifies an expression

2. In talking about "experiences," perceptual or otherwise, I am following standard practice. But it is unclear whether, ordinarily, we have visual (auditory, etc.) *experiences*, where this connotes the having of, or the undergoing of, emotional or attitudinal episodes of some sort. If I see a chair, I do not have an experience of a chair; I do not have an emotional or attitudinal reaction to the chair. Rather, I simply *see* the chair, minus *any* sort of emotional or attitudinal reaction. For now, let me point out that my usage of "visual experience," "auditory experience," etc. should not be taken to imply that simple seeings, hearings, etc. are mental episodes that we have or undergo.

referring to a creature. In general terms, "S (creature) is conscious" means that S is awake or sentient as opposed to comatose or in a dreamless sleep. Observe that, in these cases, the creature that "conscious" denotes is non-relationally conscious, i.e., conscious *sans phrase*. Philosophers, following Malcolm (Armstrong and Malcolm 1984, p. 3), sometimes refer to this type of consciousness as *intransitive* consciousness, or, following Rosenthal (1990), as *creature* consciousness. Following Carruthers (2000), I will refer to it as *intransitive creature* consciousness.

"Conscious" is often used, as well, with the particle "of" plus a noun phrase or with the connective "that" plus some sentence, as in: "Smith was conscious of the tree," or "She was conscious that he was being modest in his demeanour." In these cases, the term "conscious" functions as a verbal adjective requiring a noun phrase modified by "of" or some sentence introduced by "that" to make it intelligible. Malcolm refers to this type of relational consciousness as *transitive* consciousness. Following Carruthers, I will refer to it as *transitive creature* consciousness. Succinctly put, if a creature is transitively consciousness, then the creature is conscious *of* something or *that* something is so.

Two preliminary points concerning transitive and intransitive creature consciousness are worth making. First, as Dretske (1993, 1995) has pointed out, it is possible for a creature to be conscious *of* a thing, F, without thereby being conscious *that* it is an F.[3] Dretske calls the former "thing awareness" and the latter "fact awareness." This distinction is roughly analogous to Sellars' distinction between sentience and sapience, the distinction between hearing, seeing, and smelling, on the one hand, and believing, knowing, and thinking on the other.[4] My auditory consciousness of, say, a piano being played is very different from my *belief* that a piano is being played. Visual, auditory, olfactory, and other perceptual or sensory states do not, arguably, require the mobilization of concepts, whereas propositional attitude states like beliefs, thoughts, and judgements surely do. I may see a tree, and thereby have a visual consciousness of a tree, without knowing *that* it is a tree I am seeing. Bats, marmosets, dogs, and elephants can see, hear, and taste a whole host of objects, presumably, but they will not be aware *that* they see, hear, and

---

3. Though it seems the converse does not hold; that is, it is not possible to be conscious *that* F without, thereby, being conscious *of* G, where G is a noun phrase created from the sentence F such that G denotes the state of affairs denoted by F. Intuitively, if I am conscious that the grass is green, then I am conscious of the grass being green.

4. See Sellars 1997, sec. 1. This distinction also echoes, though certainly does not parallel, Brentano's distinction between being merely aware of an object, i.e., having a *presentation*, and making a judgement, taking a stance, or otherwise having an attitude toward that object. See Brentano 1874, pp. 78–9.

taste these objects.[5] As non-discursive animals, they do not have a mastery of the relevant concepts. A mouse, to use Dretske's example, can smell burning toast, and thus be aware of burning toast, but it will not be aware that toast is burning. These examples illustrate the difference between "experiences of *k* (as F) and thoughts about *k* (that it is F)" (Dretske 1995, p. 10).[6]

Second, Malcolm maintains that although transitive creature consciousness implies intransitive creature consciousness, the converse does not hold. That is, it is possible, according to Malcolm, for a creature to be intransitively conscious without being concurrently transitively conscious, i.e., without being conscious of something or that something is so. He invites us to consider a person who has regained consciousness after being knocked out, arguing that we can say that the person is conscious without needing to add an "of" or a "that." However, although Malcolm is unquestionably right that transitive creature consciousness implies intransitive creature consciousness – one cannot be conscious of something unless one is awake[7] – there is reason to think that

---

5. It would be rash, however, to suppose that these animals do not know what these objects *are*. It is typical of those impressed with the work of Sellars to insist that to know what an object is it is necessary to respond to that object "under a description." A creature, on this view, can know what sort of thing X is only if the creature is "able to link the concept of Xness up with other concepts in such a way as to be able to justify claims about X's" (Rorty 1979, p. 183). To have knowledge of what X is, then, it is necessary to have the concept X, which presupposes a grasp of other concepts. Thus, since infants and animals respond to objects merely discriminatively, they do not know what these objects are. This view seems unduly dismissive of the conceptual powers of at least some animals, however. A dog will not be aware that it is a Frisbee he is chasing, since the dog does not have the requisite language, i.e., does not have the concept *Frisbee*. But there seems to be no compelling reason to doubt that the dog has certain rudimentary concept-exercising abilities by which he knows what a Frisbee is.

6. I should note that Dretske notes the distinction between experiences of *k* (as F) and thoughts about *k* (that it is F), but he also notes the distinction between experiences of *k* (as F) and simply experiences of *k*. According to Dretske, seeing (hearing, etc.) an object is, in its most basic form, "non-epistemic," in the sense that one can see an object and yet have no beliefs about it (see Dretske 1969, ch. II). While all seeing *that* involves beliefs (and therefore concepts), seeing *k*, which is an F, does not imply seeing that *k* is an F, or seeing that *k* is anything else. Thus, on Dretske's view, one can have a visual experience of a tree without having a visual experience of the tree *as a tree* or, indeed, as of anything. To be sure, Dretske does not deny that for S to see *k*, *k* must look some way to S: "I am quite willing to admit that, in a certain sense, [*k*] must look some way to S (*not* to be read as: must look *like* something to S) in order for S to see [*k*]" (Dretske 1969, p. 9). However, he denies that it follows from this that, in seeing *k*, S must have some kind of belief about *k*.

7. Except, of course, when one is dreaming. Dreaming is a type of consciousness, in my view, a way of being conscious of whatever one dreams. In addition, some tricky cases suggest themselves. For example, we can imagine someone who, while sleeping, moves her leg in reaction to

he is mistaken in denying the converse implication, i.e., that one cannot be conscious without being conscious of something. In the case he describes, it seems plausible to suppose that the person's consciousness, rather than lacking content entirely, has content that is indeterminate.[8] I use the term "content" here in its most general sense to refer to anything, however vague, of which a creature is conscious, e.g., a rainbow, a hallucinatory rainbow, butterflies in the stomach, or a general feeling of depression.[9] All consciousness, as Merleau-Ponty has observed (1962, p. 5), is consciousness of something, but this "something" is not necessarily an identifiable object, let alone an object under some description.[10] The locution "S (creature) is conscious" is idiomatically *useful*, in that it gives us a way to express the indeterminacy of the content of S's consciousness. But there are no grounds here for supposing that S's consciousness lacks content altogether. On my view, then, if a subject S is intransitively conscious, then there is an *x* (or any number of *x's*) such that S is transitively conscious of *x*.

A compelling reason for thinking that all consciousness is transitive is that there seems to be no difference between a conscious state that lacks content entirely and an unconscious state. It seems a condition on the very possibility of being a conscious state that the state has content. On this view, having content is at least partly constitutive of the very nature of a conscious state, determining what it is.

This is closely related to a point for which I argue later (ch. 4), namely, that all conscious states have *intentionality*. Intentionality, in philosophical parlance,

---

a pinprick, or someone who is caring for a baby, falls asleep, and later awakens upon hearing the baby cry. A distinction must be recognized, however, between unconscious awareness and conscious awareness. Reacting to a stimulus while asleep, even if it causes one to wake up, is not, arguably, a case of consciousness at all, but rather a case of unconscious awareness.

8. Also see Armstrong's reply to Malcolm in Armstrong and Malcolm (1984, p. 117), and Honderich (1988, p. 79).

9. It is worth noting the traditional distinction between *content* and *object*. (See Woodruff Smith 1989, pp. 7–10, for a concise review of some considerations that force this distinction.) Every conscious experience has content, but some experiences, such as hallucinations, do not have an object. Nevertheless, we should like to say that even hallucinatory experiences are experiences "of" something, i.e., that they have content. If I hallucinate a pink elephant, there is no object of my experience, but the content of my experience is something that appears to me as pink and elephant-like.

10. It might be objected that, lest we acknowledge the existence of strange mental particulars or hoary old "sense-data," hallucinatory consciousness is not consciousness of anything. But this way of approaching the issue is, I think, misguided. When one hallucinates (or dreams about, or imagines), say, a pink elephant, one is indeed conscious of something, namely, something that is pink and elephant-like. But it is by no means obvious that accepting this view commits one to the existence of sense data.

is, very roughly, a property of "being about" an object that may or may not exist. According to Brentano, who appropriated the term from the Scholastic philosophers of the Middle Ages and reintroduced it into philosophy, intentionality is a distinguishing feature of all mental phenomena.[11] Paradigmatically, beliefs, hopes, fears, and desires are intentional mental states, inasmuch as each has an intentional object, an object it is about or directed at. If one is afraid or hopeful, one is always afraid *of something* or is hopeful *for something*. Perceptual states, as well, are usually regarded as having intentionality, since they are always directed at an object(s). If one hears, for example, one always hears *something*, however vague or indistinct. We often ask people *what* they hear (or what they are listening to), but it would be absurd to ask someone whether or not they are simply hearing. And in response to the question, "What do you hear?" it makes no sense to respond by saying: "I hear, but I do not hear anything."

The view that all conscious states are intentional, however, runs strongly counter to current philosophical orthodoxy.[12] It is often maintained that although some conscious states are intentional, there are other states, especially certain feelings and bodily sensations, that do not take an object and therefore possess no intentionality. Pains, itches, and tickles, for example, do not seem to be about or directed at anything. There are also certain feelings of depression and other moods that do not seem to exhibit intentionality. I shall argue, however (ch. 4), that we should dispense with the distinction between intentional and non-intentional mental states and maintain, along with Brentano, that *all* mental phenomena, at least all conscious mental phenomena,[13] have intentional content. The explanatory benefits of adopting this approach to consciousness will become clearer later on.

2.3 State consciousness

Thus far I have been chiefly concerned with the locution "S (creature) is conscious" and the two-place predicates "S (creature) is conscious of *x*" and "S (creature) is

---

11. Brentano writes: "Every mental phenomenon is characterized by what the Scholastics of the Middle Ages called the intentional (or mental) inexistence of an object, and what we might call, though not wholly unambiguously, reference to a content, direction toward an object (which is not to be understood here as meaning a thing), or immanent objectivity" (1874, p. 88).

12. Some noteworthy exceptions are Tye (1995), whose views on the matter I discuss in Chapter Four, Byrne (2001), and Crane (2003).

13. Brentano held that all mental phenomena are conscious, eschewing the now-orthodox philosophical belief in the existence of unconscious mental states. In fact, he used the term "consciousness" as synonymous with "mental phenomenon" or "mental act" (1874, p. 102).

conscious that *y*." However, philosophers are keen to distinguish these locutions from locutions signifying the existence of what they have called "state consciousness," i.e., locutions in which "conscious" is used as an adjective modifying reference to some mental state, as in: "Her thought that Paris is in France is conscious," or "She had a conscious dislike for his arrogant demeanour." In these sentences, consciousness is being attributed, not to a creature, but to a state, event, or process in (or of) that creature. Since mental states cannot be conscious of something or that something is, state consciousness is always intransitive.

Now my concern in this book is primarily with state consciousness; that is, I am concerned with the question of what constitutes a conscious state. And I will shortly offer two directs argument for (ch. 4), and subsequently proceed on, the intuitive assumption, shared by many, that a mental state is conscious if and only if the subject is, in some suitable way, aware of being in it.[14] However, a potential misunderstanding should be cautioned against. The literature on consciousness, especially the literature on higher-orderism, is rife with talk of the "consciousness of states" (Rosenthal 1990, p. 731); of "forms of consciousness which are properties of mental states" (Carruthers 2000, p. 13); of "whether a particular mental state is conscious" (Güzeldere 1995, p. 790); of "the fact that some mental states are conscious" (Van Gulick 1995, p. 271); and of the assumption "that mental states are conscious – that they possess. . .*state consciousness*" (Lurz 2003, p. 23). However, none of these authors has declared that he is to be understood as employing "state consciousness" as a shorthand expression for something like, "a state a creature is in." On the contrary, state consciousness is alleged to be a substantive type of consciousness, quite distinct from any type of creature consciousness.

But this is doubly problematic. First, supporters of the notion of state consciousness frequently advert to what might be called "semantic ascent" (see Quine 1960), that is, they advert to "the way we speak" (see, e.g., Rosenthal 1993a, p. 355; 1990, p. 730; Carruthers 2000, p. 13), in supposing that consciousness can be a property of mental states. But the ordinary use of the phrase "conscious mental state" does not licence this inference. In ordinary usage, the phrase signifies either a mental state *of which* one is conscious or a mental state one is in, *while conscious*; and neither of these uses warrants the conclusion that mental states themselves can be conscious. In short, then, the ascription of consciousness to mental states,

14. I will also offer arguments suggesting that a mental state is conscious if and only if there is something *it is like* for the subject to be in it (see ch. 3). My strategy, ultimately, is to show that, appropriately construed, the self-awareness involved in conscious states is *one and the same* as the what-it-is-likeness involved in conscious states.

though sanctioned by grammatical custom, should not be taken to indicate that state consciousness does not imply creature consciousness.[15]

It might be objected that state consciousness need not imply creature consciousness because "the dreaming subject is not conscious (hence requiring that wakefulness be necessary and sufficient for creature-consciousness), although the dreamer may be undergoing mental states which are conscious" (Carruthers 2000, p. 10). I am more disposed, myself, to say that dreams are mental states of which the dreamer is aware, from which it follows that state consciousness and creature consciousness do not come apart in dreams. After all, most dreaming subjects are capable of giving (sometimes very detailed) reports of at least some of their dreams, and it is difficult to see how this is possible if they are in no way aware of them.

Another reason to doubt that state consciousness is distinct from any type of creature consciousness is that mental states, in and of themselves, are neither conscious nor unconscious, in any sense (though it is acceptable to *talk* of conscious or unconscious mental states). To suppose otherwise is to breach of what Bennett and Hacker call the *mereological principle*, the principle that psychological predicates apply, logically, only to whole creatures, and not to their constituent parts (properties, states, etc.).[16] To be sure, some predicates *can* apply to both a whole and to its parts. For example, there is nothing logically awry about applying the predicate "has cirrhosis" to both S and to S's liver; and it makes perfectly good sense to apply the predicate "has gangrene" to both S and to S's foot. But these cases are quite beside the point. The range of predicates under consideration – viz., psychological predicates – have no intelligible application to parts (properties, states, etc.) of a whole, i.e., to parts of a conscious creature (cf. Cowley 1991, p. 66).[17] Bennett and Hacker explicate the principle as follows:

---

15. Though it would be unwise to conclude, at least without further argumentation, that creature consciousness implies state consciousness, since there are examples suggesting that the former can occur without the latter. Milner and Goodale (1995) have argued, for example, that the human brain contains two functionally distinct visual systems, one in the parietal lobes, concerned with the on-line guidance of movement (the "how" system), and one in the temporal lobes, concerned with conceptual thought and planning (the "what" system). However, evidence from *visuo-spatial neglect* patients who have damaged "what" systems but intact "how" systems suggests that the parietal system functions *automatically*, outside of conscious awareness. Yet some might be inclined to say that someone whose actions were appropriately guided by the details of his circumstances (e.g., grasping a moving object in mid-air, with hand and finger grip appropriate to the dimensions and orientation of that object) was thereby conscious (creature conscious) of those circumstances without occupying a conscious state.

16. Also see Kenny (1984), who argues that ascribing psychological predicates to anything but the whole organism is to commit what he calls the "homunculus fallacy."

17. Apparent examples of the ascription of psychological predicates to parts of a creature abound in the philosophical literature. S. White, for example, claims that "many of the sophisticated

> Our point is a conceptual one . . . . Psychological predicates are predicates that apply essentially to the whole living animal, not to its parts. It is not the eye (let alone the brain)[18] that sees, but we see with our eyes (and we do not see with our brains, although without a brain functioning normally in respect of the visual system, we would not see) . . . .The organs of an animal are parts of the animal, and psychological predicates are ascribable to the whole animal, not to its constituent parts.
>
> (2003, pp. 72–3)

The crucial point is that when one ascribes a psychological predicate to a constituent part of a creature, one is not saying something false, for to say something false is to say something that makes sense. Rather, when one ascribes a psychological predicate to a constituent part of a creature, one is saying something that, as Kant

---

subsystems of a human subject – that is, many of the more sophisticated homunculi that make up such a subject – are themselves conscious" (1987, p. 149); Searle claims that "of course our brains can think" (1984, p. 36); McGinn claims that "supervenience assures that the brain has *some* property which confers consciousness upon it" (1991c, pp. 204–5); and Chalmers suggests that "a physical system such as a brain [can] also be an experiencer" (1996, p. xi). My impression is that, in these cases (with the exception, perhaps, of White), talk of consciousness as being a property of the brain or of some other part of a creature is a convenient verbal habit or an indication of expository incautiousness. Others, though, seem sincerely to believe that the brain has psychological properties and capacities. Consider, for example, the following quote from Dennett, in which he ponders what happens in the brain when one walks into a room that is decorated with wallpaper made up of an array of identical photographs of Marilyn Monroe: "Now, is it possible that the brain takes one of its high-resolution foveal views of Marilyn and reproduces it, as if by photocopying, across an internal mapping of the expanse of wall? That is the only way the high-resolution details you used to identify Marilyn could 'get into the background' at all, since parafoveal vision is not sharp enough to provide it by itself. I suppose it is possible in principle, but the brain almost certainly does not go to the trouble of doing *that* filling in! Having identified a single Marilyn, and having received no information to the effect that the other blobs are not Marilyns, it jumps to the conclusion that the rest are Marilyns, and labels the whole region 'more Marilyns' without any further rendering of Marilyn at all" (1991, p. 355). And consider the following remark by Churchland: "[W]e can now understand how the infant brain slowly develops a framework of concepts with which to comprehend the world. And we can see how the matured brain deploys that framework almost instantaneously: to recognize similarities, to grasp analogies, and to anticipate both the immediate and the distant future" (1995, p. 4). So, according to Dennett, the brain can (or cannot) "go to the trouble" of doing something; it can "identify" things; it can "jump to conclusions"; and it can "label" things. And, according to Churchland, the brain can "develop a framework of concepts," "recognize similarities," "grasp analogies," and "anticipate the future." But, of course, the brain cannot do any of these things; only conscious creatures can.

18. The brain is not the only subject of such misascription. Descartes, for example, regularly ascribed psychological predicates to the mind, not realizing that conscious creatures does not have minds (or bodies for that matter) and that, even if we did, it would make no sense, save metaphorically, to ascribe psychological predicates to them.

might put it, "transgresses the bounds of sense." Thus, the mereological principle is a logical principle: it cannot be invalidated by empirical considerations. It could be changed by stipulation, but not, as Bennett and Hacker aver, "without changing a great deal else, [and] thereby altogether changing the meanings of our words and the structure of a multitude of familiar concepts. For the principle that psychological concepts apply to the animal as a whole and cannot be applied to its parts is held in place by a ramifying network of conceptual connections" (2003, p. 81).[19]

Bennett and Hacker invoke behavioural criteria as the primary ground for ascribing psychological predicates to creatures, arguing that these criterial grounds "are partly constitutive of the *meaning* of the predicate" (2003, p. 83). Under normal circumstances, if a person smiles profusely, shouts out "Yippee!", and jumps up and down, we take such behaviour to be justifying grounds for ascribing a state of delight or happiness to the person; and these behavioural criteria are part of what the predicates "is delighted," "is happy," etc. mean.[20]

The point here is that if one accepts the mereological principle, as one should, then the *literal* ascription of consciousness to a mental state or experience is meaningless. Consciousness is a property of persons – or, more generally, creatures – and it is the person who may become conscious of a certain experience that he is undergoing. If, for example, I realize that I am afraid, and so become aware of my occurrent feeling of fear, it is not my fear that has become conscious, but rather *I* who have become conscious of my fear.

The use of phrases like "conscious mental state" and "conscious state" need not cause confusion, however. It is reasonable to suppose that subjects or creatures are conscious of things, i.e., transitively conscious, precisely in virtue of being in certain states, so one need only stipulate that one's use of "conscious mental state" and its cognates in discussions of transitive creature consciousness is meant as a kind of *façon de parler*, i.e., as a convenient way of speaking of a range of states that make it possible for *us* to be conscious of things. Moreover, it is admissible to speak of intransitive state consciousness as long as we recognize that a mental state that is intransitively conscious is not a mental state that has the property of being conscious.

---

19. Bennett and Hacker extract from the work of Wittgenstein the insight that psychological predicates can legitimately be applied only to whole creatures. Wittgenstein expresses this insight most perspicuously with the following terse remark: "only of a human being and what resembles (behaves like) a living human being can one say: it has sensations; it sees, is blind; hears, is deaf; is conscious or unconscious" (1953, §281).

20. Bennett and Hacker are not espousing any kind of behaviourism. Their claim is that behavioural criteria are *partly*, not wholly, constitutive of what psychological predicates mean.

Rather, a mental state that is intransitively conscious is either a mental state of which the *subject* is conscious or a mental state a creature is in *while conscious*.

2.4 Unconscious mental states

It is customary to suppose that it must be possible for mental states to be unconscious.[21] In fact, for a significant number of philosophers (and scientists) it is something of an inviolable assumption that "consciousness" and "mentality" are not coextensive, that there is a wide range of contentful yet non-introspectible unconscious cognitive operations "lying behind" our conscious states.

The evidence is inconclusive, however. One of Rosenthal's favourite examples involves unconscious thinking: "We often consciously puzzle over a question about what to do or how to solve a problem, only to have the answer occur to us later, without the matter having in the meantime been in any way consciously before our mind" (2002c, p. 407; cf. Carruthers 2000, p. 19). According to Rosenthal, even "though it doesn't seem, from a first-person point of view, that we were thinking about the issue, it's clear that we must have been" (ibid.). But this conclusion seems rash. Why not simply conclude that the solution to the puzzle occurred to one after not having thought about it for a while? Unconscious headaches are another oft-cited example. It is possible to have a headache, it is claimed, that one is aware of intermittently; and during the periods when one is not aware of it, the headache persists as an unconscious headache. Again, however, this conclusion seems rash. Why not conclude that it is possible to be *peripherally* aware of one's headache and *focally* aware of something else? There are periods when a subject may not *attend* to his headache, but why suppose that during these periods the subject is not inattentively aware of his headache? Most commentators have failed to notice (or at least have neglected to explicitly acknowledge) that if unconscious pains are possible, then it is possible for one to undergo (violent?) pains of which one is entirely unaware. But it is difficult to make sense of the claim that it is possible for one to be entirely unaware that one is in pain. (When one takes pain alleviating medication, does the medication remove the pain, or does it make the pain unconscious? Is there a difference?) In any event, since the view I defend in this book does not demand that I take a position either way, I will not to delve into this issue too deeply. My view, as it happens, is that the notion of an unconscious pain – i.e., a pain of which the subject is entirely unaware – just like the notion of

21. Throughout this book I will use the term "unconscious" as synonymous with "non-conscious," so my usage here (and elsewhere) of "unconscious" is not meant in any Freudian sense.

LIVERPOOL JOHN MOORES UNIVERSITY
LEARNING SERVICES

an unconscious tickle or an unconscious itch, is incoherent. However, it seems possible to have unconscious beliefs, where an unconscious belief is taken to be "something I believe, believing which colours my emotional reactions and informs my actions, but which I am *unwilling to acknowledge*, either to myself or to others, *as* something which I believe" (Bennett and Hacker 2003, p. 269). Additionally, there does not seem to anything *prima facie* objectionable about the notion of an unconscious desire, construed along roughly similar lines.

2.5 Self-consciousness

A further notion, that of self-consciousness, requires mention, if only to distinguish its various senses from the sense with which I am principally concerned. In ordinary usage, as Ryle has documented (1949, pp. 156–7), the term "self-conscious" typically connotes embarrassment or paying heed to one's character. The term is also used, less so in ordinary usage and more so in technical usage, as a synonym for "introspection," which refers to our ability to deliberately shift our attention "inward" and think about or reflect on our own mental states. The ability to introspect is an extremely sophisticated cognitive ability possessed by only a handful of species. Introspection is, if you will, *explicit* self-consciousness, whereby we attend, either casually or attentively, to our own mental states, or to ourselves having those mental states, or simply to ourselves. I can, for example, think about my occurrent feeling of love for someone else, in which case I am explicitly conscious of one of my psychological states. But I can also think of myself as being in love, in which case I am thinking not about my love per se, but rather of myself as being in love. Moreover, I can think simply of myself, i.e., of my "global" state, so to speak. I might, for example, in a moment of reflection, think that I am a generally well-rounded person. In each of these cases, I am thinking about myself or of certain of my psychological states. Note that, in introspection, I am *thinking* about myself or of certain of my psychological states. As Searle has pointed out (1992, p. 144; cf. Shoemaker 1996b and Bennett and Hacker 2003), introspection should not be modeled on vision, where there is a clear distinction between the object seen and our visual consciousness of the object. During an act of introspection one cannot literally step back and "look" at one's mental states. Thus, introspection ought to be construed, innocently enough, as our ability to think, either casually or attentively, about our own mental states, an ability the possession of which by human beings at least seems scarcely controvertible.

Another sense of "self-consciousness" differs markedly from the foregoing senses, namely, the Cartesian sense according to which self-consciousness involves consciousness of an immaterial, enduring self. Influenced by Hume, however, few,

if any, philosophers these days hold that self-consciousness is consciousness of a "self" in any sense of the word Descartes espoused. Hume observed that introspection does not reveal an underlying self; when we introspect, all we find is a bundle of perceptions.[22]

Hume's critique of the Cartesian conception of self-consciousness is compelling, but the rather weak or low-level type of self-consciousness at issue in this book is not vulnerable to the sort of objection Hume directs at the Cartesian conception. For a creature to be self-conscious (or self-aware) in this weak sense is just for it to be aware, in a sense to be explained in Chapter Six, of its own mental states, and for it to be capable of distinguishing between itself and other things. Even non-human animals, arguably, are capable of this kind of self-consciousness. More, of course, needs to be said about this type of self-consciousness, but for now I will only point out that to say that a creature is self-conscious in this minimal sense is to say that there is *something it is like* for the creature to have the conscious experience it is having.

2.6 Phenomenal consciousness

I take, and will hereafter use, the expression "phenomenal," as it is relevant to the topic of consciousness, to mean, roughly, "appertaining to what a conscious state *is like* from the first-person perspective." Hence, if there is something it is like for a creature to be in a particular conscious state, then the creature is in a phenomenally conscious state. Paradigm examples of phenomenally conscious states are pains and sensations. There is something it is like, for example, to experience a pounding migraine, a cramp in one's foot, or a throbbing toothache. Perceptual states, as well, are unproblematically phenomenal. There is something it is like to gaze at a sunset, smell a rose, walk barefoot on a beach, and hear the clear whistles of a meadowlark. Importantly, the phrase "what it is *like*" is not to be understood as signifying a comparison, so that our conscious states *resemble* something. Rather, as Nagel has noted, the phrase means "*how* [a conscious state or experience] is for the subject himself" (1974, p. 440n6, my italics).

Following Nagel, furthermore, it is reasonable to suppose that phenomenally conscious states are not confined to discursive or linguistically endowed creatures. Although it is unclear where to draw the line as we descend the scale of phylogenetic

22. Hume writes: "For my part, when I enter most intimately into what I call *myself*, I always stumble on some particular perception or other, of heat or cold, light or shade, love or hatred, pain or pleasure. I never catch *myself* at any time without a perception, and never can observe any thing but the perception" (1978, bk. I, ch. iv, §6).

complexity, it is safe to assume that human infants and other non-linguistic creatures, even those with relatively primitive nervous systems, have conscious experience.[23]

But what is the relation between being in a conscious state and being in phenomenal state? Nowadays it is widely supposed that the two are *contingently* related, in the sense that, at least conceptually speaking, one can be in a conscious state and not be in a phenomenal state. However, as I have indicated, I want to contend that the two are *non-contingently* or *conceptually* related, in the sense that a conscious state could not – could not possibly – lack phenomenal character and exist.[24] In others words, all and only conscious states are phenomenal states. My seeing the tree, on this account, is a case of conscious seeing just in case there is something it is like for me to see the tree.

Before defending this claim, however, it will be useful to take a closer look at the notion of phenomenal character, a task that will occupy me for the remainder of this chapter. Above I characterized phenomenal character in terms of Nagel's "what-it-is-like" locution, but there are compelling dismissals of the locution in the philosophical literature that its users have yet to refute. Lycan, for example, contends that the phrase is multiply ambiguous, going as far as to claim that it is "worse than useless: it is positively pernicious and harmful, because nothing whatever is clarified or explained by reference to it" (1996, p. 77). Hacker argues, similarly, that "it is misconceived to suppose that one can circumscribe, let alone define, conscious experience in terms of there being something which it is like for a subject to have it" (2002, p. 166).[25] There is good reason to suppose, however, that there is no better idiom in the English language to convey what the essential feature of a conscious state is. To demonstrate the truth of this claim, I will provide an extended indirect argument that seeks to defuse a number of arguments for the

---

23. For some, to be sure, this is contentious. Descartes, for example, is reputed to have held that all animals are unconscious automata, and Harrison (1991) and Carruthers (1989, 2000) have recently defended versions of this Cartesian doctrine. In my judgement, however, there are no plausible reasons for conceding phenomenally conscious states to human beings and denying them to animals.

24. To this it might be objected that Quine's (1951) critique of analyticity shows that the notion of a conceptual truth is incoherent. But does it? To my mind, the proposition expressed by "It is false that a bachelor is an unmarried male" is false by virtue of the linguistic conventions governing "bachelor" and "unmarried male." I don't see how else it *could* it be false. Surely it is not false for the same reason as the proposition expressed by "Ravens are white" is false; that is, it is not false because all observed bachelors have been unmarried males.

25. Also see Malcolm 1984, pp. 45–66; Hofstadter and Dennett 1981, pp. 403–14; Hanfling 2001, pp. 48–52; and Levi 1997, pp. 461–2.

truth of its negation. To this end, I will examine Hacker's denigration of the phrase and, more briefly, Lycan's. Focusing on their criticisms will allow me to broach other pertinent issues along the way.

2.6.1 Hacker and Lycan's criticisms of the "what-it-is-like" locution

Some sketchy remarks on the notion of *qualia* will facilitate clarification of Lycan and Hacker's opposition to the "what-it-is-like" phrase, since this notion figures prominently in both their critiques. A definitional inexactness, to put it mildly, has always beleaguered this term and its singular "quale," but even when philosophers differ in their usage of the term, they seem to be attempting to capture the same concept. The following are some definitions of qualia, which are fairly representative of the current literature:

> Qualia, if there are such, are properties of sensations and perceptual states, namely the properties that give them their qualitative or phenomenal character – those that determine 'what it is like' to have them.
>
> (Shoemaker 1996a, p. 121)

> [A] mental state is conscious if there is something it is like to be in that mental state. To put it another way, we can say that a mental state is conscious if it has a qualitative feel – an associated quality of experience. These qualitative feels are also known as phenomenal qualities, or qualia for short.
>
> (Chalmers 1996, p. 4)

> Qualia include the ways it feels to see, hear and smell, the way it feels to have a pain; more generally, what it's like to have mental states.
>
> (Block 1994, p. 514)

> Qualia are the properties of phenomenal states that determine their qualitative character, i.e. 'what it is like' to have them.
>
> (Gennaro 1996, p. 7)

> [Q]ualia in sense modality M (for S) are the way objects phenomenally appear or seem to S in M.
>
> (Dretske 1995, p. 73)

> "Qualia" is an unfamiliar term for something that could not be more familiar to each of us: the way things seem to us.
>
> (Dennett 1988, p. 619)

All these authors have very different views on qualia, but a common thread running through their respective definitions is that qualia are the purely qualitative features of experience, i.e., the way things *seem* or *appear* to an experiencing subject. So defined, the notion of qualia is conspicuously similar to the notion of phenomenal character.

In fact, "qualia" is frequently used as a synonym for the phenomenal or what-it-is-like features of experience, as evidenced by Shoemaker, Chalmers, Block, and Gennaro's definitions. On this usage, "qualia," "phenomenal character," "qualitative character," and "what-it-is-like" are all just notational variants of the same concept. And few are prepared to deny that conscious states have qualia in this sense. Even Dennett, an unswerving opponent of qualia, agrees that there are "*ways things seem to us*" and that we experience things as being "one way rather than another" (Dennett 1988, p. 619).

A subtle ambiguity should be guarded against, however. Dennett and Dretske's definitions of qualia portray qualia as the way *things* seem to us, whereas Shoemaker, Chalmers, Block, and Gennaro's definitions portray qualia as the way *experiences* or *conscious states* seem to us. As a result, it seems that, according to Dennett and Dretske, qualia denote what things are like, whereas according to Shoemaker and the rest, qualia denote what experiences are like. Needless to say, these definitions have divergent entailments. One's response to the question "What was the tomato like?" would differ from one's response to the question "What was your visual consciousness of the tomato like?"[26] Answering the former would involve providing a description of the properties of the tomato, whereas answering the latter would involve providing a description of the properties of one's visual consciousness of the tomato. In essence, Dennett and Dretske seem to displace qualia by making them properties of objects, not conscious states. Presumably, however, Dennett and Dretske are not suggesting that qualia are non-mental properties of perceived objects. When Dennett wrote "Quining Qualia" in 1988, qualia were typically, if not universally, thought to be mental, and Dennett based his analysis of qualia on this then-dominant conception. And although Dretske (1995) argues at length that the qualitative properties of a given perceptual state are exhausted by the perceived object's properties, he bases his analysis of qualia on the original supposition that qualia are properties of conscious states. I take it, then, that phrases like "the way things seem to us" and "the way things appear to us" are roughly elliptical for "the way seeing (or hearing, or tasting, or whatever) an object is like for us."[27]

---

26. Cf. Carruthers, who draws an important conceptual distinction between "what the *world* (or the state of the organism's own body) is like for an organism, and what the organism's *experience of the world* (or of its own body) is like for the organism" (2000, pp. 127–8). He calls the former *worldly subjectivity* and the latter *experiential subjectivity*.

27. Perhaps I am being too charitable. Dennett and Dretske do use the phrase "the way things seem to us," or some such construction, to characterize qualia, and it is unclear whether they mean "the way seeing, hearing, etc. is like to us." But I am not sure what else they could mean by "the way things seem to us." Things can, of course, seem a certain way to us, but in describing the way something seems, one provides a description of the object one saw, heard, etc., not a description of one's conscious experience of that object (cf. Hacker 1987, ch. 6).

Talk of qualia, at least in the guise I have presented it, seems rather innocuous, so what is Dennett attempting to "quine" when he attempts to quine qualia?[28] Dennett maintains that the notion of qualia should be discarded not because he is uncertain that our conscious experiences have experiential content; rather, he inveighs against qualia because of doubts about certain other features they supposedly possess, viz., ineffability, essential privacy, intrinsicality, and noninferential and incorrigible accessibility in the first person. For Dennett, conscious states are like something, surely, but they have none of these dubious features. For this reason, Dennett, likening his strategy to that of an exasperated person abandoning a hopelessly snarled kite string, recommends that we simply eliminate the notion of qualia altogether.

For purposes of the present discussion, however, we can prescind from questions about whether qualia possess the questionable features itemized above and say, harmlessly, that qualia constitute the way seeing something, hearing something, tasting something etc. is like for us. This is roughly how Lycan defines a quale in the strict sense, though, importantly, he does not use "quale" as a synonym for "what-it-is-like":

> A quale is the introspectible monadic qualitative property of what seems to be a phenomenal individual, such as the color of what Russell called a visual sense datum. For example, if S is visually healthy and looking at a ripe tomato in good light, the tomato will look red to S, and if S focuses her introspective attention on the corresponding subregion of her visual field, S will see that subregion as an individual red patch having a roundish shape. The redness of that phenomenal patch is the quale of the containing visual sensation.
>
> (1996, pp. 69–70)

It might be supposed that Lycan's claim that a quale is an *introspectible* monadic property suggests that he thinks qualia have that dubious epistemic dimension of being directly and incorrigibly apprehensible to consciousness, but, again, we can abstract away from this potentially contentious aspect of Lycan's definition.

On, then, to Lycan's criticisms of the "what-it-is-like" locution. According to Lycan, the phrase is "ambiguous, as between phenomenal character, i.e., a quale in the strict sense, and the conscious experience of such a quale, or rather what one knows in virtue of having such an experience" (1996, p. 77). Lycan later expands on this criticism, saying that we must distinguish

> between *merely* registering a quale. . .and consciously doing so, i.e., being aware of it. Recall. . .the ambiguity of "what it's like," as between phenomenal character

28. See Dennett 1988. Dennett takes the verb "to quine" from his 1987, a satirical dictionary of eponyms. It means to deny resolutely the existence of a seemingly obvious and undeniable phenomenon.

> (i.e., a quale in our strict sense) and the conscious experience of such a quale or the introspective *knowledge of* phenomenal character. *Remember, it is important, both for psychologists and for philosophers, to separate questions about qualia from questions about awareness and introspective consciousness.*
>
> (1996, p. 100)

So, on Lycan's view, there is a clear difference between merely having an experience, i.e., registering a quale in the strict sense, and having a *conscious* experience. Lycan identifies having a conscious experience with being introspectively aware of having that experience: "consciousness is a perceptionlike second-order representing of our own psychological states and events" (1996, p. 13). Consequently, a mental state is conscious, according to Lycan, just in case one is introspectively aware of the mental state itself.

It would seem, then, that Lycan's dissatisfaction with the "what-it-is-like" locution stems from his overall theory of consciousness. Lycan identifies being conscious with being introspectively aware of one's own mental states, so merely registering a quale will not suffice for consciousness, since one might not be introspectively aware of the quale one registers. In short, one can be in a qualitative state – i.e., be registering a quale – without, thereby, being in a conscious state. The "what-it-is-like" locution, then, is ambiguous as to its intended target: does it pertain to merely registering a quale or being introspectively aware of the quale one registers?

On the view I wish to advocate, however, the supposed ambiguity of which Lycan speaks is a non-issue. Specifically, according to my preferred construal of Nagel's what-it-is-like formula, what-it-is-likeness or phenomenal character is a criterion for state consciousness, providing its necessary and sufficient conditions. Thus, a creature S is in a conscious mental state just in case there is something it is like for S to be in that mental state. This means that if S registers a quale, then S consciously registers a quale just in case there is something it is like for S to register that quale. Similarly, if S is introspectively aware of a quale, then S is consciously introspectively aware of that quale just in case there is something it is like for S to be introspectively aware of that quale. On this proposal, "what-it-is-like" is decidedly unambiguous, since questions about being in a conscious state cannot be prized apart from questions about what it is like for the subject to be in that state.

Let me go on to remark parenthetically that given the controversy surrounding the terms "qualia" and "quale," I will henceforward abandon these terms altogether, except if they are required to expound someone else's view. As noted previously, "qualia" is sometimes used relatively harmlessly to mean the what-it-is-like features of experience, but, in addition to the term's association with controversial metaphysical notions like incorrigible accessibility, it is also used sometimes to denote atomistic *feels*, as Lycan's use of the phrases "registering a quale" and "visual sensation" suggests. But, to state a point I enlarge on below (and in subsequent chapters), the very notion of a "qualitative feel" (or of "registering a quale" or of a "visual

sensation"), as this notion applies to perceptual (especially visual) consciousness, is, despite its perennial appeal, fundamentally wrongheaded.

I turn now to consideration of Hacker's critique of the "what-it-is-like" locution, which is significantly more sustained than Lycan's. Hacker aims a substantial part of his critique at the traditional notion of qualia, taking it more or less for granted that "qualia" is synonymous with "what-it-is-like." As Hacker sees it:

> An experience is held to be a *conscious* experience just in case there is something which it is like for the subject of the experience to have it. Consciousness, thus conceived, is defined in terms of *the qualitative feel of the experience*. This qualitative feel, unique to every distinguishable experience, is alleged to be *what it is like for the subject of the experience to have the experience*.
> (2002, p. 162)

Observe that Hacker defines conscious experience – or, more accurately, takes others as defining conscious experience – in terms of the qualitative *feel* of the experience. Hacker proceeds to criticize this view, contending that:

> if we were to ask a normal person how it feels to see the table, chair, desk, carpet, etc., etc., he would wonder what we were after. There is nothing distinctive about seeing these mundane objects. Of course, seeing the table differs from seeing the chair, desk, carpet, etc., but the difference does not consist in the fact that seeing the desk *feels different* from seeing the chair. Seeing an ordinary table or chair does not evoke *any* emotional or attitudinal reaction whatsoever in normal circumstances. The experiences differ in so far as their objects differ.
> (2002, p. 162)

Now I concur with, and invite the reader's concurrence with, every word in this latter passage, the central constituent idea of which is clear and persuasive. But notice that here Hacker's objection to the idea of defining perceptual experience in terms of what-it-is-like is predicated on the supposition that "what-it-is-like" is synonymous with "qualitative feel," as the first of the quoted passages indicates. However, although it is customary to define perceptual experience in terms of qualitative feels, nothing about the expression "what-it-is-like" makes associating it with these putative feels compulsory. In other words, one can agree with Hacker that it is erroneous to say perceptual experiences have distinctive *feels*, and yet affirm that there is something *it is like* for the subject of a perceptual experience to have it. Otherwise put, "qualitative feel" and "what-it-is-like" are by no means unavoidably coextensive.[29] It follows that Hacker, despite correctly assailing the claim

---

29. The tendency to equate what-it-is-likeness and qualitative feel is certainly not due to Nagel. Nagel never even mentions "feels" in his 1974. For Nagel, what-it-is-likeness is essentially bound up with "subjective character," but for an experience to have subjective character is not necessarily for it to have a "feel."

that atomistic feels distinguish our visual experiences, is too hasty in denying that seeing a chair, table, desk, etc. is like something.

Now Hacker might concede the foregoing arguments and still maintain that any attempt to explicate conscious states in terms of what-it-is-likeness will founder. After propounding the objections just surveyed, he proceeds to attack the very grammar of the "what-it-is-like" locution, arguing that the "very expression 'There is something it is like for a person to have it' is malconstructed" (2002, p. 167).[30] In addressing Hacker's criticisms it will be useful to look at the grammatical contexts in which the "what-it-is-like" locution, in discussions of consciousness, typically occurs. It is most often found in the following contexts:

(1) "There is something it is like for an S to be an S," where S is some creature;
(2) "There is something it is like for S to V," where V is a psychological verb or verb phrase;
(3) "There is something it is like for S to be in a state of V-ing"; and
(4) "There is something it is like for S to be in M," where M is a noun phrase denoting some mental state.

Since (3) and (4) are nominalizations of (2), (2) – (4) are cognates, and, for the most part, I have defended the locution as it occurs in these contexts. But Nagel's original version of the what-it-is-like formula contains an occurrence of (1). For Nagel, a creature (Nagel uses "organism") has conscious mental states just in case there is something it is like *for the creature to be that creature*. According to Hacker, however, the expression "There is something it is like for an S to be an S" is objectionable because it illegitimately reiterates the subject term in the object position. One can reasonably ask "What is it *like for a woman* (as opposed to a man) to be a soldier?", but the question "What is it like for a soldier to be a soldier?" is, Hacker contends, almost certainly awry. Consequently, it makes no sense to ask "What is it like for a bat to be a bat?", since the question is vulnerable to the same objection of illegitimate reiteration. We can interpret the question as equivalent to "What is it like to be a bat?", but this question is an inquiry into the characteristic attitudinal features of the life of a bat, not an inquiry into what constitutes a bat's being in a conscious state (Hacker 2002, pp. 167–70).

I am inclined to agree with Hacker that the expression "There is something it is like for an S to be an S" is problematic, so I propose to steer clear of the expression altogether. More fully, I propose to use the "what-it-is-like" locution exclusively

30. Others (e.g., Lormand 2004, Hellie 2007b) have offered careful, sometimes painstaking examinations of the grammatical structure of the "what-it-is-like" locution, but, as far as I know, only Hacker has arrived at the conclusion that it is wholly ungrammatical.

as it occurs in (2) – (4) above, i.e., as it occurs in contexts such as "There is something it is like for S to V." Hacker is perhaps justified in expressing unease about the expression ". . . for an S to be an S,"[31] but the expression ". . . for S to V" and its cognates are, I want to suggest, perfectly licit and informative.

Now Hacker is equally averse to the expression ". . . for S to V," so it will be useful to consider his objections to it. According to Hacker, "What we cannot say is:

(1) 'There is something which it is like to V.'

let alone

(2) 'There is something it is like for [S] to V'" (2002, p. 166).

(1), Hacker argues, "is apt only for cases of comparison," whereas (2) "is a miscegenous [sic] crossing of the form of a judgment of similarity with the form of a request for an affective attitudinal characterization of an experience" (2002, p. 166).

One cannot but agree with Hacker that (1) is apt for cases of comparison, at least at first glance. And, again at first glance, it seems that (2) is vulnerable to the charge that Hacker levels against it. The "at first glance" qualification is important, for – as Hacker acknowledges and as Nagel pointed out upon advancing the what-it-is-like formula – the "what-it-is-like" locution is not intended to indicate a comparison of any kind; so it is a matter of some perplexity that Hacker objects to the locution on the grounds that it goes awry unless it is employed in cases of comparison. In accordance with his late Wittgensteinian sympathies, Hacker is keen to emphasize, not unwisely, that ordinary language cannot be abused without generating conceptual confusion. But even Hacker accepts that it is sometimes necessary to posit technical concepts or replace ordinary concepts with technical ones for specialized purposes, which is clearly Nagel's intention in employing the "what-it-is-like" locution in a non-comparative sense. Nonetheless, Hacker's basic strategy is to argue, on ordinary language grounds, that the philosophical use of the "what-it-is-like" locution should be discarded. Specifically, he argues that the philosophically uninitiated do not use (1) and (2) the same way philosophers do, i.e., they do not use (1) and (2) except to elicit a comparison, a judgement of

---

31. My agreement here with Hacker is guarded. Hacker argues that the question "What is it like for an S to be an S?" can be interpreted as "What is it like to be an S?", but, as I mentioned, he then dismisses this on the grounds that this question is an inquiry into the characteristic attitudinal features of the life of S, not an inquiry into what constitutes S's being in a conscious state. But I confess that the intended contrast is not clear to me. As an anonymous reviewer for my 2005 pointed out, why not say that what it is like for an S to be an S (i.e., what it is like to be an S) is the sum of all the things it is like for S to V, for a range of salient and relevant values of V?

similarity, or an emotional or attitudinal reaction. And, further, it is difficult to get a grasp of what the philosophical use of (1) and (2) *means*.

To understand what the philosophical use of (1) and (2) means, it is important to recognize that conscious states are never sub-personal, which means that (1) is an abbreviated version of (2). That is, the expression "There is something it is like to V" is always shorthand for "There is something it is like for S to V," where S is some creature (cf. Lormand 2004, pp. 307–8). When a bat, for example, avails itself of the properties of sonar to locate prey, there is something it is like *for the bat* to avail itself of the properties of sonar. Similarly, when a human being sees a tree, there is something it is like for the human being to see the tree. In fine, there is always something it is like for the subject of a conscious experience to have that experience.

Admittedly, the expression "There is something it is like for the subject of the experience to have it" may seem obscure at first, especially if cast in non-comparative terms, but the expression is not, as Hacker would have us believe, nonsensical. In his 1974, Nagel tried to explain the meaning of the expression in terms of the difference between the subjective and the objective. But in an earlier essay he (perhaps unwittingly) explained its meaning more crudely, yet I think more illuminatingly, as a "type of internality" (1965, p. 353). Notice that when a creature consciously perceives an object (consciously thinks about an object, has a conscious sensation, etc.), there is something "going on upstairs," something going on that is entirely lacking in non-conscious artifacts and organisms, i.e., a type of internality. And I can think of no better idiom in the English language that expresses this fact about consciousness more fittingly than the idiom of "what-it-is-like." When a creature consciously perceives *x*, consciously thinks about *x*, etc., the creature undergoes a type of internality, i.e., there is something[32] it is like for the creature to perceive *x*, consciously think about *x*, etc.

32. The "something" here functions, obviously, as a variable, as Lormand (2004, p. 308) has pointed out. I offer an account of what this variable generalizes over, i.e., an account of *what* it is like for a creature to perceive *x*, think about *x*, etc. in Chapter Eight (§8.2).

CHAPTER 3

A formula for state consciousness

3.1 Nagel's what-it-is-like formula

In the preceding chapter I indicated how we might understand the "what-it-is-like" locution. This result puts us in a favourable position to examine Nagel's what-it-is-like formula and assess its viability against some alleged counterexamples. The formula, as Nagel presents it, runs as follows:

(N1) "An organism has conscious mental states if and only if there is something that it is like to *be* that organism – something it is like *for* the organism."[1]

(1974, p. 436)

As it stands, N1 is ambiguous, so in the course of educing from it a substantive formula for state consciousness, I will need to make a number of adjustments and refinements to it, ultimately delivering a formula quite different from anything Nagel likely intended.

I have already intimated one of my main reasons for rejecting the what-it-is-like formula as Nagel presents it. Recall that, in the previous chapter, I proposed eschewing the problematic expression "There is something it is like for an S to be an S" in favour of "There is something it is like for S to V." But Nagel presents his version of the formula using the former expression, arguing that there is something it is like for a bat to be a bat. In the following revisions to the formula, then, I have replaced the occurrence of "...for an S to be an S" with ". . .for or an S to V," or some analogue thereof.

Consider, then, the following version of the formula:

(N2) A creature S consciously perceives *x* at *t* if and only if there is something it is like for S to perceive *x* at *t*.

It is unnecessary, for my purposes, to take a stance on whether N2 is true, since this reading of Nagel's formula says that a *creature* consciously perceives *x* at *t* (i.e., is transitively conscious of *x* at *t*) just in case there is something it is like for

1. The phraseology here is conspicuously Kantian. In the Transcendental Deduction of the *Critique* Kant repeatedly emphasizes that consciousness is, if anything, a first-personal phenomenon: it is something *to us* or *for us* (see Kant 1965, A120, B132, B138).

the creature to perceive *x* at *t*. N2, then, pertains to *transitive creature* consciousness, not to *intransitive state* consciousness. Replacing the relevant occurrence of "creature" with "mental state" and temporally indexing the occurrence of S's mental state to correspond to the occurrence of S's what-it-is-like state remedies this lack:

(N3) A mental state M of a subject S is conscious at *t* if and only if there is something it is like for S to be in M at *t*.

Stated in terms of perceptual consciousness, the Nagelian formula for state consciousness is essentially this:

(N4) A perceptual state P of a subject S is conscious at *t* if and only if there is something it is like for S to be in P at *t*.

When I see the tree, I am in a conscious visual state only insofar as there is something it is like for me to see the tree.

(N4) extends in a natural way to different types of non-perceptual conscious states, such as conscious emotions:

(N5) An emotional state E of a subject S is conscious at *t* if and only if there is something it is like for S to be in E at *t*.

And we can adjust the formula to accommodate a state of conscious thinking:

(N6) A cognitive state C of a subject S is conscious at *t* if and only if there is something it is like for S to be in C at *t*.

It is worth emphasizing an important feature of this proposed criterion for a conscious state. On the proposed view, occurrent states are conscious if and only if there is something it is like for the subject to have them. Take, for example, the conscious activity of solving a difficult mathematical problem. I can solve a difficult mathematical problem and later reflect on what the activity was like. But my reflecting on my solving the problem is not what makes the actual activity of solving the problem a conscious activity. In order for the solving of the problem to be a conscious activity, there must be something it is like for me – *while* solving – to solve the problem.

3.2 Putative counterexamples

Although most philosophers are favourably disposed to the claim that at least some conscious states have phenomenal character, my claim that a creature is in a conscious state if and only if there is something it is like for the creature to be in that state is considerably more controversial. Several philosophers have developed arguments and adduced cases purporting to show that although phenomenally

conscious states undoubtedly exist, it does not follow that we are committed to the claim that a creature is in a conscious state just in case there is something it is like for the creature to be in that state. We should, they urge, emend the biconditional and say, more innocuously, that a creature is in a conscious state *if* there is something it is like for the creature to be in that state. On this view, state consciousness and phenomenality are dissociable: a mental state's being like something for a creature is sufficient, but not necessary, to make that state conscious.[2]

Block (1995) has argued, for example, that a distinction must be recognized between phenomenal consciousness and access consciousness, both being cases of full-fledged consciousness. According to Block, phenomenal consciousness and access consciousness are easily and dangerously conflated, and there are, he argues, conceptually possible instances of one without the other.[3] For example, well-documented cases suggest that it is possible for subjects with blindsight to respond to visual stimuli in the absence of visual awareness. Blindsight patients have typically suffered brain lesions in certain portions of the striate cortex (the brain's "visual" centre) that cause substantial areas of blindness, or scotomata, within their visual fields. Yet, with the appropriate methods of testing, they are able to reach for, point to, or issue remarkably accurate statements about visual stimuli presented within their areas of blindness. Paradoxically, however, although better-than-chance responses to visual stimuli can be demonstrated, subjects with blindsight invariably deny being able to "see" them. In blindsight, then, it seems that certain perceptual or discriminative capacities are preserved, while the phenomenal features typically associated with such capacities are not. According to Block, these perceptual capacities (on a specific construal elaborated below) are access conscious without being phenomenally conscious. In addition, Carruthers maintains that it is "plain that we do need some notion of access-consciousness in addition to a notion of phenomenal consciousness, because at least some states can – in a fairly intuitive sense – be conscious without there being anything which it is like to undergo them. Consider acts of thinking" (2000, p. 17). Nelkin has argued, in a similar vein, that phenomenality is not necessary for conscious thought. Nelkin asks, tendentiously: "But consider, what is it like to have a thought? We are all tempted to say we know

---

2. This use of "make" here, and elsewhere in similar contexts, is perhaps misleading, calling to mind some sort of causal relation. No such implication is intended. The relation between what-it-is-likeness and state consciousness is better understood as a *constitutive* relation: being in a what-it-is-like state constitutes one's being in a conscious state.

3. Since nothing at this point hangs on the notion of access consciousness, I momentarily ignore the details of what Block means by it.

LIVERPOOL JOHN MOORES UNIVERSITY
LEARNING SERVICES

exactly what it's like to have a particular thought, say, the thought of needing to get the pistons in one's car checked. But what is it that we know?" (1987, p. 235).

Armstrong's well-known example of the daydreaming, long-distance truck driver is occasionally held up as problematic for the biconditional as well. "After driving for long periods of time," Armstrong says, "particularly at night, it is possible to 'come to' and realize that for some time past one has been driving without being aware what one was doing. The coming-to is an alarming experience. It is natural to describe what went on before one came to by saying that during that time one lacked consciousness" (1981, p. 723). Armstrong goes on to deny that one lacks consciousness during such phases of inattentiveness, but he argues that although one has minimal and perceptual consciousness, one lacks state consciousness. By minimal consciousness Armstrong means that one is not asleep or in a coma; by perceptual consciousness he means that one is conscious of stop signs, curves in the road, and so on. As it see it, then, Armstrong's example, under its intended interpretation, is designed to show that transitive creature consciousness does not imply state consciousness, i.e., that it is possible for a creature to be conscious of something without being in a conscious state. So construed, the example does not bear against the what-it-is-like formula, since the formula purports to provide necessary and sufficient conditions for a conscious state (i.e., the formula says that a *state* of a creature is conscious at *t* if and only if there is something it is like for the creature to be in that state at *t*); it does not purport to provide necessary and sufficient conditions for transitive creature consciousness (i.e., the formula does not say that a *creature* consciously perceives *x* at *t* if and only if there is something it is like for the creature to perceive *x* at *t*). Nonetheless, one might try using the example to attack the what-it-is-like formula; that is, one might argue that, during the phase of inattentiveness, the driver is in a *conscious perceptual state* that lacks phenomenal character.

The above cases, however, are not deleterious to the what-it-is-like formula.[4] I will examine them in turn, starting with Block's increasingly popular distinction between access consciousness and phenomenal consciousness.[5]

---

4. There are doubtless other arguments purporting to demonstrate the possibility of conscious states that lack phenomenal character, but the ones I expound on here are the most oft-cited and widely considered to be the most persuasive. Thus, I am reasonably confident my objections possess some general applicability.

5. Among others who accept Block's distinction between phenomenal consciousness and access consciousness are Burge 1997, Byrne 1997, Chalmers 1996, and Rosenthal 2002c.

3.2.1 Access consciousness and phenomenal consciousness

Block's usage of "phenomenal consciousness" (P-consciousness) is the same as my usage of "phenomenal character." Block takes P-conscious states to be states of consciousness that are like something. In contrast, access consciousness (A-consciousness) is a functional notion that has to do with reasoning and information processing. A state is A-conscious, on Block's view, if its content is freely available for direct control of reasoning and action (1995, p. 382). Paradigm examples of A-conscious states are beliefs, thoughts, and desires. These are propositional attitude states, which are characteristically expressed by means of a psychological verbal expression followed by a "that" (or cognate) clause.

Some remarks, however perfunctory, on the notion of representation will be useful at this point, since Block employs this notion in his explanation of A-consciousness. Philosophers, psychologists, cognitive scientists, etc. seem to know what they are talking about when they employ terms like "mental representation," "representational content," "representational state," and so on, but the sheer existence of mental representations (hereafter simply "representations") is eminently disputable: when I see *x* (hear *x*, etc.), *x* is *pre*sented to me,[6] but it is not *re*presented to me. To be sure, if "representation" is being used in a stipulative or otherwise technical sense, then it may be apposite to say that there are at least visual representations. Neuroscience has established inductive correlations between seeing various properties in one's environment, such as blueness and greenness, and certain patterns of neural activity in the striate cortex. These correlations have enabled us to reliably infer from the fact that a specific neural event is occurring in S's brain at *t* that S sees a specific property in her visual field at *t*; and they have also enabled us to map properties in the visual field – e.g., the contours of a visible object – on to specific cell excitations in the striate cortex (see Bennett and Hacker 2003, pp. 305, 307). Thus, if we get rid of the ordinary meaning of "representation," and denominate the former representations of the latter, then it would be appropriate to say that certain neural firings in the striate cortex represent orientation, depth, etc. However, I can find no stipulation in the writings of Block and other representationalists to the effect that they are using "representation" in

6. Though it is unwise to say, as philosophers often do, that my *experience* presents *x* to me. If A presents something *x* to B, then, roughly, A brings forth, introduces, renders, shows, or exhibits *x* to B. But my visual experience of, say, a barn does not bring forth, introduce, render, show, or exhibit the barn to me. Rather, if the barn is in my visual field and perceptual conditions are normal, I will see it, and thereby have a visual experience of it. But in having a visual experience of the barn, my experience doesn't present the barn to me.

a specialized sense, so it is entirely unclear what they mean when they claim that we have visual representations. For if "representation" is used in accordance with its ordinary use, then it makes no sense to say that certain neural firings represent orientation, depth, etc.; for, in its ordinary sense, "representation" means "to stand for," "indicate," or "symbolize," and neural firings do not stand far, indicate, or symbolize anything.

All is not lost, however. The term "representation," at least as it is used in contemporary philosophical psychology and cognitive science, is, surely, a term of art, much as the term "idea" as used by the Empiricists was a term of art. (The lack of a widely accepted definition, even a rough one, of the term "representation" bears out this observation.) It is therefore to be expected that a definitive definition of the term is unavailable and probably not forthcoming. My best guess is that a representation, at least on Block's understanding, is an informational entity, endowed with causal powers, that is instantiated in the brain. Very crudely, representations occur when external stimuli "impinge" on sensory systems that are capable of detecting, receiving, and manipulating such stimuli. If a creature encounters a red ball, for example, the redness and the circularity (among other things) of the ball might impinge on the creature's visual system. The creature's visual system, designed (evolutionarily or in the lab) to detect these properties, processes the information received and generates a representation of a red ball. This representation may, or may not, reach the system that is in control of rational action and reasoning. This very rough characterization is silent, however, on whether a representation is a mental image of some sort or a purely syntactic entity like a string of symbols.

Now in the course of explicating A-conscious states, Block says that "A-conscious content is representational," and that "it is the essence of A-conscious content to play a role in reasoning, and only representational content can figure in reasoning" (1995, p. 383). He goes on to say that "it is in virtue of its representational content, or the representational aspect of its content, that a state is A-conscious" (ibid.). And later he explains that "what makes a state A-conscious is what a representation of its content does in a system . . . . Content is A-conscious in virtue of (a representation with that content) reaching the Executive System, the system that is in charge of rational control of action and speech" (1995, pp. 383–4). Block is here claiming, in effect, that propositional attitudes and other access content integrally and essentially involve representational content. And since only representational content can figure in reasoning, states with the appropriate representational content are A-conscious when they reach the executive system and play a role in verbal report, rational action, and reasoning. Thus, a state of a creature is A-conscious if the creature adjusts her behaviour rationally, makes a verbal report, or makes a rational judgement in light of some representational content possessed by that state (and perhaps does so cleverly, inventively, etc.).

According to Block, one way of showing that it is conceptually possible for there to be A-consciousness without P-consciousness, is to show that it is conceptually possible for there to be a creature with beliefs, thoughts, and desires without there being anything it is like for the creature to have such mental states. "If there could be a full-fledged phenomenal zombie," Block maintains, "say a robot computationally identical to a person, but whose silicon brain did not support P-consciousness, that would do the trick" (1995, p. 385).[7] Block, however, is keen to resist such far-fetched examples,[8] so he invites us to imagine a less controversial – though still hypothetical – case, viz., a case of what he calls super-blindsight. Real blindsight patients can guess accurately only when given a choice among a very limited, usually binary, set of alternatives. But suppose, Block says, "that a blindsight patient could be trained to prompt himself at will, guessing what is in the blind field without being told to guess. The super-blindsighter spontaneously says 'Now I know that there is a horizontal line in my blind field even though I don't actually see it'" (1995, p. 385). On Block's view, this is a clear, conceptually possible instance of a conscious state that lacks phenomenal character. The super-blindsighter is in a state in virtue of which

---

7. Despite its outlandishness, a number of philosophers take the notion of a phenomenal zombie very seriously, as evidenced by the spate of recent literature on the subject. See, among countless others, Balog 1999; Brueckner 2001; Carruthers 2000; Chalmers 1996; Levine 2001, esp. ch. 2; and Worley 2003. The recent resurgence of interest in zombies is largely due to Chalmers, who advances a modal conceivability argument to the effect that since zombies are conceivable, zombies are logically possible. And since zombies are logically possible, reductive physicalism is false. This argument goes back to at least Descartes (Sixth Meditation), who argued that we can conceive of all the physical and functional facts obtaining in the absence of consciousness. But Chalmers thinks he can dodge what Brueckner calls the "standard objection" to the argument – according to which conceivability is not a guarantor of logical possibility – through his *two-dimensional semantics* (see Chalmers 1996, ch. 2). Most commentators do not think Chalmers' version of the conceivability arguments works, though Worley (2003) argues that it is not vulnerable to the standard objection.

8. And wisely so, for compromising the example right from the start is the question of whether we should even ascribe mental states to robots. Most of us are disinclined to accept that a robot, even a very sophisticated modern robot, has desires and beliefs. Consider the computational state within the robot's "brain" that is designed to represent low battery power. Should we identify this computational state with the robot's desire for more battery power? Or consider the computational state within the robot's brain that represents light intensity. Should we identify this computational state with the robot's belief that there is light in its visual field? It seems that we should no more say that the robot believes that there is light in its visual field than we should say that a thermometer believes that the temperature it is gauging is hot or cold. To be sure, the robot undergoes more sophisticated representational activities than the thermometer, but it does not follow that the robot has desires and beliefs.

he is conscious, in some sense, of the horizontal line, yet there is nothing it is like for him to be in such a state. Of course, the super-blindsighter has a thought that there is a horizontal line in his blind field of which he is both A-conscious and P-conscious, but it is not the thought with which Block is concerned. Rather, he is concerned with the state of the super-blindsighter's perceptual system that gives rise to the thought. It is this state, according to Block, that is A-conscious without being P-conscious. The super-blindsighter is in a conscious perceptual state, but his perceptual state lacks phenomenal character.

Now Block explicitly acknowledges that he has a "liberal terminological policy" (1995, p. 390), but I want to suggest that there is something amiss about employing the term "conscious" to describe the state of the super-blindsighter's perceptual system that gives rise to his thought that there is a horizontal line in his blind field. One consideration that lends a degree of initial plausibility to this claim is that ascribing A-consciousness to this state of the super-blindsighter's perceptual system commits one to ascribing A-conscious states to artefacts and organisms that, by all accounts, do not have conscious states. Consider Cog, a sophisticated robot at MIT that possesses a whole host of self-protective, self-regulatory, and self-exploring mechanisms.[9] Cog has impressive cognitive talents, yet presumably it does have conscious states. To be sure, Cog has various internal states in virtue of which it is able to perform various tasks, such as detect an object in its path and manoeuvre around it. But these are not conscious states. More plausibly, Cog, to appropriate some terminology from Flanagan (1992), has internal states that are *informationally sensitive* to its environment, but not *experientially sensitive*: Cog possesses non-conscious computational-informational states that control its behaviour. Arguably, Block's super-blindsighter possesses a similar non-conscious informational state that enables him to react to the horizontal line.[10] It is reasonable to suppose that many simpler organisms, like bees, ticks, and tubeworms, possess these kinds of non-conscious informational states as well. These organisms, in virtue of occupying such states, react to environmental stimuli, i.e., they gather (or passively receive) information from the environment and use it to control their behaviour. The point to be drawn from these considerations is that if, on the basis of his informational sensitivity to the horizontal line, we attribute an A-conscious state (whose content is the horizontal line) to the super-blindsighter, then it seems

9. Dennett discusses Cog in his 1998, ch. 9.

10. To forestall misunderstanding, it should be stressed that I am not saying that the super-blindsighter's cognitive reaction to the horizontal line – i.e., his thought that there is a line in his blind field – is unconscious; rather, I am saying his *perceptual* reaction is unconscious.

that we cannot, at least not without inconsistency, deny various A-conscious states to Cog (whose content is whatever objects or properties. Cog is informationally sensing), since Cog possesses the same kind of informational sensitivity as the super-blindsighter. But Cog does not have conscious states;[11] hence, the super-blindsighter does not have the relevant A-conscious state.[12]

One might be tempted to respond that this is all too quick, that we can tell a story about A-conscious states in terms of the complexity of the responses to objects, facts, properties, etc. in light of the representational content of the system in question. And once we have reached a certain level of complexity, we can say that a system has conscious states even if we are uncertain about whether the system has phenomenal states. I am not sure whether proponents of a conceptual distinction between state consciousness and phenomenal character would want to pursue this line of argument, however, since it is by no means obvious that we have not already developed non-conscious systems – Cog, for example – that undergo representational activity that is almost as sophisticated as the representational activity involved in the super-blindsighter's perception of the horizontal line (see Dennett 1998, pp. 160-8). Moreover, the objection assumes that it is possible to state, in a non-arbitrary way, at which point the representational activity of the system in question becomes sufficiently complex to warrant affirming that the system has conscious states. Is the representational activity sufficiently complex when it results in responsiveness to a representation that engenders behaviour that mimics human behaviour? Probably not, since certain non-conscious organisms and artefacts – Cog, for example – exhibit very humanlike behaviour as a result of purely non-conscious representational activity. Is the representational activity sufficiently complex when the system in which it occurs can use it as a tool in learning? Again, probably not, since evidence suggests that consciousness is not even necessary for certain kinds of learning (see, e.g., Jaynes 1976, pp. 31–6). And if consciousness is not even necessary for certain kinds of learning, it is difficult to see why learning should be indicative of state consciousness. Complexity will not do.[13]

---

11. This observation applies only to the current state of technology. Though I am sceptical about the likelihood that we will ever devise ways for silicon brains to support consciousness, there is nothing incoherent in maintaining that the development of such technology is, at least in principle, possible. In addition, we do not want to preclude the possibility that consciousness has evolved elsewhere in the universe in systems that are not carbon-based. These are empirical matters, however, not to be settled *a priori*.

12. Nagel (1974, pp. 436–7) makes a similar point.

13. This conclusion is apparently endorsed by McGinn: "saying that 'complexity' is the missing ingredient [for consciousness] is blatant handwaving, no kind of theory at all" (1995, p. 246).

One might also object that there is a key difference between Cog and the super-blindsighter that my argument has overlooked, namely, that the former, unlike the latter, is not (creature) conscious at all, and thus not capable of having conscious states. According to this objection, only creatures that we deem conscious in the first place can be said to have conscious states, and since we do not deem Cog conscious, it cannot be said to have conscious states. This objection therefore takes aim at the first premise of my argument – i.e., that if the super-blindsighter's perception of the horizontal line is A-conscious, then Cog has A-conscious states – and says that even if the super-blindsighter's perception of the horizontal line is A-conscious, it does not follow that we are committed to ascribing A-conscious states to Cog.

But there is a major problem here. Block's rationale for ascribing A-consciousness to the super-blindsighter's perception of the horizontal line has nothing to do with the super-blindsighter's being a creature that we deem conscious. On the contrary, Block's view, as we have seen, is that the super-blindsighter's perceptual state is A-conscious in virtue of its functional role, i.e., in virtue of its potential impact on the Executive System. But notice that, by this criterion, Cog's "perceptions" of objects in its environment are also A-conscious, since they have a potential impact on the Executive System. The point is that Block's argument could proceed on the assumption that the super-blindsighter is an entirely non-conscious system, a phenomenal zombie. But if the argument could proceed on that assumption, it is quite beside the point that the super-blindsighter happens to be a creature that we deem conscious.

The objections I have just rebutted to my approach to the case of the super-blindsighter assume that Cog cannot plausibly be described as a creature that has conscious states, but another possible concern about my approach is that it could be taken as verifying what many have long suspected, namely, that the concept of state consciousness is simply not amenable to analysis in terms of necessary and sufficient conditions, but is, rather, a variable concept, not subsumable by any single principle and requiring modification as different cases present themselves. I have argued, by modus tollens, that if we ascribe an A-conscious state (whose content is the horizontal line) to the super-blindsighter, then we seem committed to ascribing A-conscious states to Cog, inasmuch as Cog possesses the same kind of informational states as the super-blindsighter. And since Cog does not have conscious states, the super-blindsighter does not have an A-conscious state whose object is the horizontal line. But my modus tollens might be another's modus ponens: an objector might insist that we affirm the antecedent of the above conditional and conclude, instead, that we need to modify our concept of state consciousness to include artefacts like Cog. After all, Cog possesses a number of self-controlling and self-exploring mechanisms and engages in what is describable as intelligent behaviour. So what prohibits us from ascribing conscious states to it?

We are debarred from ascribing conscious states to artefacts like Cog, at least at this stage of technological development, because Cog's computational-informational

states are not *first-person knowable*. In fact, Cog cannot be said to *know* anything about its computational-informational states. There are a number of complicated questions surrounding the nature of first-person knowledge, e.g., how it is acquired and how it differs from third-person knowledge. But, as Kriegel has pointed out (2004, pp. 197–8; cf. Gennaro 1996, p. 119), we need not answer these questions to hold that, intuitively (although crudely), a state is conscious just in case it is first-person knowable. All that is required is that there *be* a distinction between first-person and third-person knowledge; and that there is such a distinction is well nigh indisputable. If I think about a blue sky, the way *I* know that I am thinking about a blue sky is very different from the way *you* know that I am thinking about a blue sky. Your knowledge that I am thinking about a blue sky is arrived at *inferentially*, whereas my knowledge that I am thinking about a blue sky is not. Put differently, your knowledge that I am thinking about a blue sky, unlike my knowledge that I am thinking about a blue sky, is *mediated*.

This observation presupposes a general principle regarding the relation between state consciousness and first-person knowledge, namely, that being first-person knowable is at least partly constitutive of what "conscious state" means. Kriegel puts it this way:

> Consider what fixes the reference of 'the conscious'. Conscious states presumably constitute a natural kind, and as such must have an underlying nature. But when we refer to a mental state as conscious, the surface feature that guides us in classifying it as conscious is what we may call the property of being *first-person knowable*. When we consider a certain phenomenon conscious, it is because our knowledge of it is first-person knowledge. Phenomena we have only third-person knowledge of are unconscious.
>
> (2004, p. 197)

This seems intuitively plausible: we would not call a mental state conscious if we had only third-person knowledge of it. We use the criterion of first-person knowledge to "home in," to use Kriegel's phrasing (2004, p. 198), on the phenomenon – i.e., conscious states – that we want to talk about.

Now if a state is conscious just in case it is first-person knowable,[14] it follows that the super-blindsighter's perception of the horizontal line is unconscious, since the super-blindsighter's perception of the horizontal line is not first-person knowable. Indeed, there is no difference, to the super-blindsighter, between his perceiving the

---

14. My (and Kriegel's presumably) use of "knowable" is not meant to indicate a "potential" of some kind. That is, I am not suggesting that a state is conscious just in case it *could* be known from the first-person perspective. Such a view would be most implausible, for it would entail, absurdly, that many unconscious mental states, since they are potentially known, are conscious. My claim is that a mental state is conscious at *t* just in case the subject of the mental state has first-person knowledge of it at *t*.

line and his not perceiving the line; his perceiving the line, quite unlike his thought that he perceives the line, occurs at an entirely sub-personal level. This means that any knowledge the super-blindsighter has of his perception of the horizontal is third-person knowledge. Just as *you* have to infer that the super-blindsighter perceives the horizontal line (from his verbal or non-verbal behaviour), so the *super-blindsighter* has to infer that he perceives the horizontal line (from his verbal or non-verbal behaviour or from his own thought).

Notice that this argument cannot be averted by maintaining that the super-blindsighter has first-person knowledge not just of the output of his perceptual state, e.g., his thought that there is a horizontal line in his blind field, but also of certain aspects of his perceptual state itself, e.g., of the fact that it is a visual state rather than, say, a tactile state. This response will not work because if the super-blindsighter has first-person knowledge that his perceptual state is a visual state, then it must be the case that there is something it is like for him to see the horizontal line; for, minimally, he is non-inferentially aware that he sees, rather than feels, the horizontal line (cf. Clark 2000, p. 30). But in what else, other than in some sort of phenomenal apprehension, could the super-blindsighter's non-inferential knowledge that he sees, rather than feels, the horizontal line consist?

Now it might be protested that, ultimately, Block is free to claim that even states that are first-person unknowable can be said to be conscious states. After all, Block admits to having a relatively lax terminological policy, so there is nothing barring him from claiming that he is simply using "conscious" more liberally. To suppose otherwise, this objection continues, is to establish the conclusion that there is a conceptual connection between state consciousness and phenomenal character by *verbal fiat*.

This objection misfires. While it is true that philosophers are free to extend the notion of a conscious state beyond its normal (admittedly sometimes vague) boundaries for specialized purposes, there is no reason to stretch it to a point at which it becomes legitimate to ascribe consciousness to states that are first-person unknowable (and thereby to the internal functional/informational states of artefacts like Cog). This might seem a terminological quibble, but it isn't. Block, in his writings on consciousness, pronounces on one aspect or another of a phenomenon denoted by the term "consciousness," but there is general agreement – in the philosophical and non-philosophical literature – regarding the meaning of this term, even if its *precise* meaning is occasionally unclear. Indeed, such general agreement is a necessary condition for a discussion of consciousness to get off the ground in the first place. But on anyone's understanding of the meaning of the term "conscious" ("consciousness," etc.), mental states that are first-person unknowable do not, and could not, count as conscious states, just as on anyone's understanding of the meaning of the term "bachelor," a married man does not, and could not,

count as a bachelor.[15] Succinctly put, to ascribe consciousness to states that are first-person unknowable is to speak not of consciousness, but of something else, perhaps consciousness*. But, of course, Block and the rest of us are interested in *consciousness*, not consciousness*.[16]

The present argument can be summarized as follows: a state is conscious just in case it is first-person knowable; the super-blindsighter's perception of the horizontal line is not first-person knowable; therefore, the super-blindsighter's perception of the horizontal line is unconscious. It follows that Block has not demonstrated that there is a conceptual distinction between state consciousness and phenomenal character, since the kinds of access states he invokes to make his case for this conclusion – e.g., the super-blindsighter's perception of the horizontal line – are not conscious states. In other words, Block has shown that there are various computational-informational states that lack phenomenal character (this seems scarcely contestable), but he has not shown that there are *conscious* states that lack phenomenal character.

It still remains to contend with Carruthers' argument that acts of thinking necessitate a need for some notion of access consciousness in addition to the notion of phenomenal consciousness because certain acts of thinking "can – in a fairly intuitive sense – be conscious without there being anything which it is like to undergo them" (2000, p. 17). According to Carruthers, phenomenal character is not necessary for conscious thought, since one can have a conscious thought without there being anything it is like to have it. This argument is devised on the presumption that if there is something it is like to have an experience, then the

---

15. To the objection that this comparison is unapt because, unlike the concept of consciousness, the concept of a bachelor does not have vague boundaries, I would refer the reader to Hanfling's musings on the matter: "Even the word 'bachelor', often regarded as an easy candidate for such definition [in terms of necessary and sufficient conditions], turns out to be more recalcitrant in this respect than one might assume. Not only must 'unmarried' be replaced by 'has never been married' so as to exclude widowers from being bachelors, but something further may need to be added to exclude celibate priests. There must also be provision to exclude men in societies where there is no concept of marriage and, finally, men in *our* society who, though not married, are living in a stable relationship with a woman" (2000, p. 95n1).

16. I cannot within the ambit of this book rebut any of the familiar (and typically hopelessly confused) objections to the argument strategy exemplified here of appealing to "what we say." Let me just point out that the strategy, which has an extraordinarily impressive philosophical lineage (Wittgenstein, Ryle, Austin, Stove, and Hanfling come to mind), is altogether undeserving of the neglect it suffers nowadays and is nearly unavoidable in a piece of exact philosophy. Expressions like "We wouldn't say. . .", "S cannot be said to. . .," "It is nonsense to say. . .," etc. pepper the writings even of those who expressly denounce the strategy.

subject of the experience is having a subjective feel of some sort. To further quote Carruthers: "[w]hile it may be true as a matter of fact that all conscious acts of thinking have subjective feel, because all such acts occur in 'inner speech' or in visual or other forms of imagery, it does not seem to be part of the very concept of a thought that this should be so" (2000, p. 17). As we have seen, however, a creature can be in a what-it-is-like state without undergoing a subjective feel of any kind, which means that Carruthers' claim that not all acts of thinking have a what-it-is-likeness is too hasty. Thus, it does not follow that certain acts of thinking are access conscious without being phenomenally conscious.

This response to Carruthers applies *mutatis mutandis* to Nelkin, since Nelkin, like Carruthers, conflates what-it-is-likeness with qualitative feel (or, at times, with having mental images) in his argument against phenomenal character being necessary for conscious thought. On Nelkin's interpretation of Nagel, "having a thought is supposed to 'feel' different from having a desire. And having this thought is supposed to 'feel' different from having that one" (1987, p. 237).[17] But, contends Nelkin, certain acts of thinking do not "feel" like anything: "Thinking that 1000-sided figures have more sides than 999-sided figures does not 'feel' different from thinking that 1001-sided figures have more sides than 1000-sided ones" (1993, p. 427). This seems right, but, for reasons given above, Nelkin is too hasty in concluding from this that certain acts of thinking lack phenomenal character.

3.2.2 The daydreaming, long-distance truck driver

Turning now to Armstrong's example of the daydreaming, long-distance truck driver, Armstrong, as will be recalled, holds that during the phase of inattentiveness the driver is in a state of minimal and perceptual consciousness (1981, p. 723). The driver is not asleep or in a coma, and he perceives stop signs, curves in the road, and so on. I have already remarked that this example, under its intended interpretation, is not designed to show that there are conscious states that lack phenomenal character. Rather, it is designed to show that creature consciousness does not imply state consciousness. However, it might be argued that we can exploit the example to show that there are conscious perceptual states that lack phenomenal character. According to this tack on the case, the driver is in a conscious perceptual state, but there is nothing it is like for him to be in that state.

---

17. Nelkin, in his 1989a, expresses his view as follows: phenomenal states (or what Nelkin calls phenomenological states) "are not essential to thinking and feeling (in the sense of emotions and affects). They seem, at most, to be essential only for states like pain sensations, visual images, auditory images, kinaesthetic sensations, and the like, i.e., those states we usually think of as sensations" (p. 131).

There is, however, a burden borne by those who want to maintain that the driver is in a conscious perceptual state during the phase of inattentiveness to provide an explanation of what constitutes the driver's being in such a state. It might be suggested that if a subject is in a perceptual state, then, *ipso facto*, the subject is in a conscious state; hence, since the driver is in a perceptual state, the driver is in a conscious state. But this line of argument is unpersuasive: as before, it may always be asked what makes the driver's perceptual state conscious. It is reasonable to assume that the driver sees the road, but in virtue of what is his seeing the road a conscious state? In virtue of what, that is, is his seeing the road a state of perceptual consciousness as opposed to a state of non-conscious, informational awareness? According to the what-it-is-like formula, a perceptual state is conscious just in case there is something it is like for the subject to be in that state. Thus, the what-it-is-like formula provides a plausible criterion for a conscious perceptual (or non-perceptual) state, whereas simply stating that perceptual states are conscious is no criterion at all. The task is to offer a criterion or set of criteria for what constitutes a conscious perceptual state. This criterion or set of criteria must spell out the role the concept of consciousness plays in the language we use to categorize perceptual states. My suspicion is that one cannot succeed in arguing persuasively for such a criterion or set of criteria without making essential reference to the what-it-is-like formula.

These considerations suffice, I think, to cast doubt on the suggestion that, during the phase of inattentiveness, the driver is in a conscious perceptual state that lacks phenomenal character. But I want to make some further remarks on the case of the long-distance truck driver, if only to provide some positive reasons for rejecting the notion that the driver's conscious state lacks phenomenal character.

The case of the long-distance truck driver calls attention to an interesting feature of our conscious experience, namely, that virtually all of our waking hours are engaged, at least to some extent, in perceptual activity. While walking to a job interview, for example, I might not pay attention to the objects I pass because of my preoccupation with how I will perform in the interview, but there is still a sense in which I perceive those objects. I hear people talking in the background; I see various objects in my way; I smell coffee brewing as I pass a coffee shop; and I feel the sidewalk beneath my feet. I perceive all these things at least to the extent that I am aware of them as objects in my environment. To describe this phenomenon more precisely, let me expand on the distinction between *focused* consciousness and *peripheral* consciousness. When I see my friend in a corridor filled with strangers, I direct my perceptual energies in such a way that focuses my attention exclusively on him. This focused consciousness is a relatively "high-grade" consciousness wherein I focus upon one item, my friend, to the exclusion of others. On the other hand, while walking down a familiar street my attention is not really

focused on the various objects I pass at all. I do not perceive these objects in a full-fledged sense; that is, I do not attend to them.[18] I do not, for the most part, direct my attention toward them as objects of my perceptual consciousness. Nevertheless, there is still a sense in which I am conscious of the objects I encounter, albeit peripherally or inattentively.

Now it might be argued in opposition to the what-it-is-like formula that although I am undeniably perceptually conscious of the objects I encounter in some weak or minimal sense, there is nothing it is like for me to be conscious of them. But is there not? If asked, I would certainly be able to provide a phenomenological description of seeing – yet not attending to – the objects I passed.[19] Such a description would involve pointing out that I perceived the objects only peripherally. It would also involve describing the experience as very unlike actually attending to an object, i.e., as very unlike perceiving an object and taking notice of it or making a judgement about it.[20] If I am looking at an object about which I make a roughly contemporaneous judgement, then my perceptual state is judgemental and perceptual; I am judgementally and perceptually attending to the object. So my description of seeing the objects I encountered would involve describing my conscious state as being unlike seeing an object and making a judgement about it. Most important, though, my description would involve pointing out that seeing the objects was, well, like seeing something, a state of consciousness that is intrinsically different from, say, smelling something, touching something, or dreaming about something.

Arguably, the long-distance truck driver's experience of inattentively perceiving his environment is similarly amenable to phenomenological description. If, after

---

18. This echoes Searle's (1992, pp. 137–9) distinction between the objects or events at the *centre* of our attention and those at the *periphery*. According to Searle, in addition to the object(s) or event(s) at the centre of our attention, there are usually a large number at the periphery, i.e., a large number of which we are only faintly aware.

19. Notice that the description at issue here is a description of my *perceptual state*, not a description of the objects I perceived. However, if one can describe the objects one perceived, then one can describe what it was like to perceive them, i.e., one can say, minimally, that seeing them (smelling them, tasting them, etc.) was like seeing something and unlike, say, hearing something.

20. An underlying supposition here is that one may go some way toward describing what a conscious state was like by describing what it was *not* like. Here is Dwyer on the matter: "There are valuable positive gains to be had from a description of our experience. . .[b]ut generally the gains from phenomenological description are negative. That is, the greatest gain from a description of experience is that it debunks various misdescriptions or fabrications of particular kinds of experience" (1990, p. 34).

"coming to," the driver is asked what his experience was like during the phase of inattentiveness, he might report that seeing the objects in his environment in this minimal sense was quite unlike seeing something and making a specific judgement about it.[21] He might also report that seeing the objects was simply like seeing something, a state of consciousness that has its own distinctive phenomenality. There are those who would have it that during the phase of inattentiveness the driver's perceptual state has no phenomenal character, but if a perceptual state is not like anything, then it should not be possible to describe, upon reflection, what that perceptual state was like. The idea here is that if a case of seeing, touching, hearing, smelling, or tasting (yet not attending to) an object is amenable to phenomenological description, then there is something it is like to be in that perceptual state.[22] To get a genuine counterexample to the what-it-is-like formula one must come up with a state of consciousness that is not like anything. But insofar as it is possible for the driver to describe what his perceptual state was like during the phase of inattentiveness, his perceptual state was like something. So stated, the relationship between our ability to describe a conscious state and there being something it is like to be in that state may seem tautological.[23]

To the foregoing argument I can anticipate an objection. To spell out this objection, I will risk being slightly tedious and recast my argument in a form more amenable to point-by-point comparison. Here, in distilled form, is the argument:

1. If the driver's perceptual state is not like anything during the phase of inattentiveness, then, after coming to, it is not possible for the driver to describe what his perceptual state was like during the phase of inattentiveness.

---

21. Though there is a puzzle here, for if the driver did not make a judgement about, say, the stop sign, how did he know to stop when he saw it? Upon seeing the sign, did the driver unconsciously make a judgement about what course of action to take? For present purposes this complication may be disregarded.

22. As Dwyer (1990, p. 34) has pointed out, it is usually impossible to provide a phenomenological description of immediate experience while one is engaged in immediate experience. One cannot describe one's experience of riding a roller coaster, for example, while riding a roller coaster. For the most part, a description of "immediate experience" is not a description of occurrent experience. The description is based on recalling what the experience was like.

23. This is not to suggest, of course, that being able to describe a conscious state is *necessary* for that state to have phenomenal character. Non-human animals cannot describe what their conscious states are like, but it does not follow that their conscious states do not have phenomenal character. I add this disclaimer to forestall the impression that I am advancing some sort of verbalizationist thesis, according to which verbalization is the essence of consciousness.

2. After coming to, it is possible for the driver to describe what his perceptual state was like during the phase of inattentiveness.
3. Therefore, the driver's perceptual state was like something during the phase of inattentiveness.

The objection concerns premise (2). This premise is far from unassailable, it might be argued, because it is not clear that the driver would not describe his experience as follows. "Seeing the obstacles I encountered wasn't like anything for me. I must have seen them, since I stopped and slowed down at the appropriate times, but I can't remember seeing them." On its face, this objection seems easily rebutted. We often forget our conscious states, making it impossible for us to describe what being in those states was like, but not remembering a conscious state does not, of course, mean that there was nothing it was like to be in that state. However, remembering a conscious state is necessary if one is to describe what that state was like. As a result, the argument will have to be amended to include an event-memory clause:

1. If the driver's perceptual state is not like anything during the phase of inattentiveness, then, after coming to, it is not possible for the driver to describe what his perceptual state was like during the phase of inattentiveness
2. If, after coming to, it is possible for the driver to remember his perceptual state during the phase of inattentiveness, then the driver can describe what his perceptual state was like during the phase of inattentiveness.
3. After coming to, it is possible for the driver to remember his perceptual state during the phase of inattentiveness.
4. Therefore, the driver's perceptual state was like something during the phase of inattentiveness.

The argument is valid, but premise (3) is vulnerable to an objection. According to Armstrong's higher-order experience theory of consciousness, one is fully conscious only insofar as one is introspectively aware of one's mental states. In the course of defending this theory, Armstrong speculates that a special connection exists between introspective consciousness and event-memory. "When the long-distance truck driver recovers introspective consciousness," Armstrong claims, "he has no memory of what happened while it was lacking" (1981, p. 727). He goes on to say:

> [I]t is tempting to suppose, therefore, as a psychological hypothesis, that unless mental activity is monitored by introspective consciousness, then it is not remembered to have occurred, or at least it is unlikely that it will be remembered. It is obvious that introspective consciousness is not sufficient for event-memory. But perhaps it is necessary, or at least generally necessary.
>
> (1981, p. 727)

Notice that to bolster his hypothesis that introspective consciousness and event-memory are intimately linked, Armstrong draws on the truck driver case. He claims that upon recovering introspective consciousness the driver has no memory of what his perceptual state was like during the phase of inattentiveness. From this he hypothesizes that event-memory and introspection are connected. It is important to be clear that Armstrong does not draw on his higher-order theory of consciousness to argue that the driver cannot remember what his perceptual state was like. He simply asserts, rather, as though it were self-evident, that, after coming to, the driver cannot remember what his perceptual state was like during the phase of inattentiveness. And since the driver cannot remember what his perceptual state was like, he certainly cannot describe what it was like; or so the argument goes. But should we accept Armstrong's claim that the driver cannot remember what his perceptual state was like?

I do not think it is necessary for a proponent of the view that conscious states and phenomenal character are not dissociable to oppose Armstrong here. For even if it did turn out that the driver is incapable of recalling his perceptual state, this would only indicate that there is no way to prove whether the driver, during the time preceding his coming to, was in a phenomenal state. We could not conclude on the basis of the driver's inability to recall his perceptual state that he was in a state of consciousness that lacked phenomenal character.

Moreover, if further empirical research suggests that people who have had such experiences are, indeed, incapable of remembering them, then we would have to entertain the possibility that such people, during the time preceding their coming to, are not perceptually conscious at all. For example, it could really be that the driver is simply too deep in thought to be conscious of the obstacles he encounters.[24] There seems to be nothing incoherent, at least, in describing the case this way, nothing prohibiting us from supposing that it is possible, in principle, that the driver was unconsciously negotiating the hazards of the road. But this tack on Armstrong's example does not, of course, show that there are states of consciousness that lack phenomenal character. It shows, rather, that we might be

---

24. Though see Searle (1992, pp. 137–9) for an opposing view. As noted earlier, Searle makes a distinction between objects at the *centre* of our attention and those at the *periphery*. According to Searle, there are, at any given time, a large number of objects of which we are only peripherally aware, but it is a mistake to suppose that we are not conscious of these objects. As Searle puts it: "When I drove to my office today, most of my attention was on philosophical thoughts. However, it is not true to say that I drove unconsciously. Unconscious driving would have led to automotive disaster" (1992, p. 138). For reasons that should be clear by now, I am inclined to agree with Searle, but, as I discuss below, we cannot preclude the possibility that we might be capable of performing certain actions unconsciously – i.e., automatically.

capable of performing certain actions automatically, i.e., in the absence of deliberate conscious control, which cannot be ruled out *a priori*. Indeed, contemporary cognitive science and psychology incline in this direction, proposing that there are various cases in which an agent is competently Φ-ing, but is not conscious of either her Φ-ing or of the objects in her surrounding environment while she is Φ-ing.[25]

Some independent corroboration for the sort of approach to the case of the long-distance truck driver defended here can be found in the work of Dennett. Dennett's approach to the case is in some respects similar to mine, and suggests, like mine, that the case does not represent a type of consciousness that lacks phenomenal character. I have argued that if, after coming to, the driver is asked to report what his perceptual state was like during the phase of inattentiveness, then he could report what his perceptual state was like, and so his perceptual state was like something. But we might suppose, along with Dennett, that the case of the long-distance truck driver is "a case of rolling consciousness with swift memory loss" (1991, p. 137). Dennett adopts this approach to the case because he thinks that if, during the time preceding his coming to, the driver had been asked about his current perceptual state, he could have provided a description of the details of at least some of his perceptions: "*surely if you had been probed* about what you had *just* seen at various moments on the drive, you would have had at least some sketchy details to report" (ibid.; cf. Byrne 2004, p. 219). So even though Dennett's approach to the case is silent about whether the driver, after coming to, can remember what his perceptual state was like during the phase of inattentiveness, it nevertheless affirms that the driver's perceptual state was like something.

As the foregoing remarks indicate, then, I am inclined to think that the driver is in a conscious mental state during the phase of inattentiveness; that is, he is in a mental state in virtue of which he is focally conscious of the contents of his daydream and peripherally conscious of the conditions of the road. But this state does not lack phenomenal character. Incidentally, this take on the case of the long-distance truck driver squares nicely with a fact about the case that is often overlooked, namely, that when the driver emerges from the state of inattentiveness, he does not say, "'Good Lord, I'm now aware of being in such-and-such sensory states'. . .[On the contrary] he says, 'Good Lord, I'm on Rte. 153 and I must have come through three traffic lights.' It is the road he realizes he has not been attending to, not his own perceptual states" (Lycan and Ryder 2003, p. 133).

---

25. I am here referring, of course, to intentional *action*, not mere behaviour, such as blinking, hiccupping, or breathing.

3.3 Non-conscious phenomenality?

I have stated the formula for state consciousness in terms of a material biconditional, but the arguments so far have predominantly gone into defending only half of that biconditional, viz., that phenomenal character is necessary for conscious states. I have yet to defend the converse implication, i.e., that phenomenal character is sufficient for conscious states. The claim that a creature can occupy a non-conscious phenomenal state may seem incredible, but the claim has, it seems, gained some currency in recent years (e.g., Gennaro 1996; Nelkin 1986, 1987, 1989a); so even if the arguments from the previous section are accepted, it might still be supposed that there is a conceptual distinction between state consciousness and phenomenal character.

I suspect, however, that the dispute here is largely terminological. It is unclear whether anyone holds that it is possible for a creature to occupy a non-conscious what-it-is-like state. Gennaro, for example, accepts the existence of non-conscious phenomenal states, but denies the existence of non-conscious qualitative states (see Gennaro 1996, pp. 7–11, 131–2). According to Gennaro, a phenomenal state can occur in the absence of its typical felt quality, but when it does it ceases to be qualitative. Gennaro, in other words, subscribes to the view that pains, itches, tickles, etc., states typically regarded as phenomenal, can occur non-consciously. But he does not maintain that it is possible for a creature to non-consciously *feel* a pain, itch, tickle, etc. The issue turns on Gennaro's use of the term "phenomenal": he does not, as I do, use the term as synonymous with "what-it-is-like." Rather, a phenomenal state, in Gennaro's terminology, is a "mental state which typically has qualitative properties" (1996, p. 121), and a "qualitative state" is a "phenomenal state with its qualitative property" (ibid.).

Similarly, Nelkin maintains that "either it makes sense to say that there are unfelt sensations or there are felt sensations (phenomenological states) we are not conscious about [sic]" (1989a, p. 129). But Nelkin is not here affirming the existence of non-conscious what-it-is-like states. To see this, consider the different types of consciousness that Nelkin distinguishes:

C1: first-order information-processing consciousness
C2: second-order, direct, noninferential accessing consciousness
CN: phenomenological consciousness
CS: sensation consciousness or image-representation consciousness.

According to Nelkin, "one is sometimes C2 about one's own C1 states" (1989a, p. 132); that is, one sometimes has direct, second-order access to one's first-order informational states. But both C1 states and C2 states, whether they occur independently or in tandem, are cognitive states, and, as we have seen, Nelkin holds that

cognitive states lack phenomenality. I have already dealt with the question of whether conscious cognitive states, such as conscious propositional attitudes, lack phenomenality, i.e., the question, in essence, of whether phenomenality is necessary for conscious thought. My answer, as will be recalled, is that there are no convincing arguments for the claim that it is not, and that those who argue for such a claim mistakenly conflate phenomenality or what-it-is-likeness with the having of "qualitative feels" (see §3.2.1 above). The problem at hand, then, is whether CN and CS states can occur in the absence of consciousness. If they can, then we would seem to have a case of non-conscious phenomenality. CN states, Nelkin tells us, are what-it-is-like states and are best treated as a subclass of CS states (1989a, p. 131); accordingly, CN states are a type of non-sentential, image-like representational state. Nelkin uses "image" here in its wider, non-visual (or not necessarily visual) sense to denote whatever representational system conscious systems actually use (1989a, p. 135). Now when Nelkin says that phenomenological or CN states can occur non-consciously, he does not, as far as I can tell, mean that it is possible for a creature to non-consciously feel a pain, or be in a non-conscious what-it-is-like state. Rather, he means that it is possible for a creature to occupy a CN state while not occupying a C2 state, i.e., that it is possible to be in a phenomenological state and not be in a concurrent second-order state directed at that phenomenological state (see Nelkin 1989a, p. 139). So, on Nelkin's view, if one construes conscious states in terms of a relation between a C2 state and some other state, then it is possible to occupy a non-conscious phenomenological state; the phenomenological state (the pain, the itch, the tickle, or whatever) will be non-conscious because it is not accompanied by a C2 state. This is what Nelkin means when he claims that it makes sense to say that there are unfelt sensations. But, for all that, it does not follow, and Nelkin does not affirm, that it is possible to be in a non-conscious felt state, i.e., a non-conscious what-it-is-like state. Once a phenomenological state becomes non-conscious, it ceases to be a felt or what-it-is-like state. There is nothing it is like for a creature to be in a non-conscious phenomenological state.

In an earlier article (1986), Nelkin appears to advance a similar argument. He maintains that pain is best treated as a cognitive state rather than as a sensation state.[26] "[T]here is," he says, "a large cognitive element in our very feeling of pain" (1986, p. 137). So, if Nelkin is right, it is possible to (1) have a pain sensation without being in pain, i.e., without occupying the appropriate affective-cognitive-behavioural state, and (2) be in pain, i.e., occupy the appropriate affective-cognitive-behavioural state, without having a pain sensation. (1), not (2), is germane to the present discussion,

26. Though this should not be taken to imply that Nelkin holds that pain is *no more than* a cognitive state. Pain, on Nelkin's view, is best treated as a cognitive state rather than as a sensation state, but Nelkin does not hold that pains lack any sensational component.

since (2) is just a variant of the claim, discussed at length in section 3.2.1, that it is possible to occupy a cognitively, yet non-phenomenally, conscious state. So (1) might be taken to imply that phenomenality is insufficient for consciousness. Let me explain. It is possible, apparently, for lobotomized patients and patients who have been given morphine after the onset of pain to feel pain yet avow that it does not hurt (here Nelkin cites Dennett 1978, esp. pp. 190–229). In addition, there are people who are said to never feel pain. Nelkin asks us to imagine that such a person has pain sensations just like those of a normal person but is not troubled by them and does nothing to mitigate them. Just like the lobotomized patients and the patients who have been given morphine, such a person has pain sensations, i.e., qualitative states, but these pain sensations occur in the absence of any appropriate affective-cognitive-behavioural state. Are these cases indicative of the existence of non-conscious qualitative states? Importantly, Nelkin never claims that they are. According to Nelkin, saying that these people are in pain after all would "only muddle things" (1986, p. 145). He prefers to treat these people, "despite their possession of our pain sensations, as not being in pain" (1986, p. 145). But he never says that their non-cognitive "pains" are non-conscious. And, indeed, even if their non-cognitive "pain" does not provoke the usual behavioural response, it seems implausible to suggest that these people are non-consciously feeling something, whatever it is.

Now Carruthers, in an effort to render somewhat palatable the outrageously counterintuitive thesis that infants and animals do not have phenomenally conscious experiences, defends the view that it is possible for creatures to non-consciously feel pains. According to Carruthers, it is "better to say that non-conscious pains would be pains that the organism feels, while denying that these states of feeling-pain would themselves have feel or any subjective 'what-it-is-likeness'" (2000, p. 201). For Carruthers, then, we are not bound by ordinary usage to equate "what-it-is-like" and "feel," nor, for that matter, are we bound by ordinary usage to equate "consciousness" and "feel." I have already had occasion to argue that "what-it-is-like" and "feel" are not strictly synonymous, since one can be in a what-it-is-like state without undergoing a feel of any kind. But the converse is the rub: can one feel something, a pain say, without there being anything it is like to feel it? Carruthers thinks so. I am inclined to think that this is something of an affront to the rules governing the use of "consciousness" and "feel," but it is not necessary to press Carruthers on this point. Notice that Carruthers does not claim that there are non-conscious what-it-is-like states. Non-conscious pains are not like anything for the creature that has them, even if the creature can "feel" them in Carruthers' technical sense.

A genuine counterexample to the claim that phenomenality or what-it-is-likeness is sufficient for consciousness is simply not in the offing. In order to demonstrate the possibility of a non-conscious phenomenal state, it is necessary to show that it is possible for a creature – in the absence of consciousness – to be in a state that there

is something it is like for the creature to be in. If it can be shown, for example, that S is non-consciously occupying pain state P (perceptual state O, etc.) despite sincerely avowing that there is something it is like to be in P, then we would have a genuine counterexample. But what could it possibly mean to say that S is non-consciously occupying a pain state that there is something it is like *for* S to be in?

3.4 Summary

The conceptual terrain in discussions of consciousness is often bumpy and sometimes barely navigable, mainly because philosophers persist in applying the word "conscious" to states that are not conscious. To be sure, some philosophers maintain that we should use the word "conscious" to refer only to states that have phenomenal character (e.g., Gennaro 1996),[27] but none, as far as I know, has made the claim to conceptual necessity advocated here. If the suggestions made in this chapter are any indication, the notion of a conscious state that lacks phenomenal character is conceptually incoherent; indeed, if I am right, it is impossible to conceive of a conscious state that lacks phenomenal character. This is a much stronger claim than has hitherto been made in the literature on consciousness.

The general idea is this. Whenever it is hypothesized that a mental state is conscious in virtue of having, say, representational content (or informational content, or whatever), we may always ask: in virtue of what is the representational content conscious, or, more specifically, what makes the mental state to which the representational content gives rise conscious? What, in other words, is the condition(s) under which we are justified in affirming that a state with that particular representational content is conscious? In contrast, we cannot coherently ask this question when it is said that a mental state is conscious because there is something it is like for the subject to have it. If it is remarked that "A subject is in a conscious mental state M because there is something it is like for the subject to be in M," it makes no sense to ask: "In virtue of what is the subject's what-it-is-like state conscious?" Intuitively, there being something it is like for the subject of M to be in M guarantees M's being conscious. Conscious states bottom out, so to speak, at what-it-is-likeness, which strongly suggests that conscious states are not dissociable from phenomenality.

To be clear, I have been concerned in this chapter with meaning of "conscious *state*," not with the meaning of "consciousness" generally. My claim is not that

27. Flanagan seems receptive to this view as well, as evidenced by his claim that "for me all and only conscious states are qualitative" (1992, p. 72).

every type of consciousness essentially has phenomenal character – e.g., there seem to be types of creature consciousness that lack state consciousness altogether and thereby lack phenomenal character (see ch. 2, n15) – but rather that *state consciousness* essentially has phenomenal character.

The claim that all state consciousness essentially has phenomenal character bears on matters of some importance to the chapters that follow. In particular, it will play a pivotal role in my argument for the claim that state consciousness necessarily involves self-awareness. I pursue this argument in the next chapter.

CHAPTER 4

Consciousness and self-awareness

4.1 Preamble

The arguments of the previous chapters were designed to converge on the thesis that a creature occupies a conscious state at a particular time in virtue of instantiating a certain property, viz., the property of there being something it is like for the creature to be in that state at that time. As it stands, this is not an *explanation* of state consciousness, but rather a reasonable constraint on the adequacy of any account of state consciousness: what-it-is-likeness or phenomenal character, since it is a fundamental and indispensable fact of conscious experience, will, perforce, have to serve as one of the basic *explanandums* in a theory of what constitutes a conscious state. The principal aim of this chapter is to begin developing a framework within which to provide an explanation of what phenomenal character *is*. To do this, I first make some remarks on intentionality, arguing that *all* conscious experiences, including pains, emotions, and moods have intentional content. I then provide a defence of the Transitivity Principle, according to which state consciousness necessarily involves self-awareness, in the sense that "a mental state's being conscious involves one's being conscious of that state in a suitable way" (Rosenthal 2000b, p. 265). I canvass some arguments for this principle, consider and then disarm an argument that purports to demonstrate that it's false, and advance two positive arguments for it. This chapter, then, is largely preparatory for subsequent chapters where, in light of the central claims of this chapter, I examine, in Chapters Five and Six, two models of consciousness that link consciousness and self-awareness – viz., the higher-order model and the Brentanian model, respectively – and, in Chapter Eight, draw a connection, a very intimate one, between self-awareness and phenomenal character.

4.2 A gloss on intentionality

Throughout the remainder of this book I will occasionally rely on a not uncontroversial claim, namely, that all conscious states have intentionality. It is important to distinguish this claim from a stronger claim, famously held by Brentano, that

intentionality is the *mark* of the mental (and therefore the mark of "the conscious," since Brentano held that all mental states are conscious), which I take to be the thesis that all and *only* mental states are intentional. The claim on which I will rely is considerably weaker, namely, that all conscious states *have* or *exhibit* intentionality. In this section I provide, with some argumentation, an explication of the gloss I will put on this claim and make some remarks on why I think the claim, though controversial, is nevertheless worth assuming in the discussions to follow.[1]

Conventional wisdom has it that perceptual states, as well as fears, desires, thoughts, etc. have intentionality. However, it is commonly supposed that a whole host of other states, such as pains, itches, tickles, and certain diffuse emotions and moods, do not have an intentional object and therefore possess no intentionality. It seems quite possible, for example, to experience a sudden sense of elation or a general feeling of depression without it being the case that there is anything these mental states are about. As Searle says: "only some, not all, mental states and events have Intentionality. Beliefs, fears, hopes, and desires are Intentional; but there are forms of nervousness, elation, and undirected anxiety that are not Intentional" (1983, p. 1). Bennett and Hacker claim, similarly, that "there are psychological phenomena that are *not* intentional at all, such as objectless moods (e.g., cheerfulness) and sensations (e.g., pain)" (2003, p. 422). And McGinn writes:

> By sensations, we shall mean bodily feelings. . .as well as perceptual experiences. They're different in an important respect, which calls for a subdivision within the class of what we are calling sensations: bodily sensations do not have an intentional object in the way that perceptual experiences do. We distinguish between a visual experience and what it is an experience of; but we do not make this distinction in respect of pains.
>
> (1982, p. 8)

While an extensive analysis of intentionality cannot be undertaken here, a few relevant points can be made. As a beginning, consider pain, a mental state commonly regarded as having no intentional content. We can, I think, cast doubt on the assumption that pains are non-directed or "objectless" by appealing to the plausible, and independently motivated, thesis that pain is, in the first place, not an *object* of consciousness, but rather *consciousness* of an object, namely, a given bodily region, usually a strained, injured, or diseased bodily region.[2] On this view, a sensation of

---

1. The claim that all consciousness is intentional is a natural corollary of – in fact is entailed by – the claim that all consciousness is transitive. Thus, it would not be unfair to characterize my discussion in this section as an extension of my discussion of transitive consciousness in Chapter Two.

2. See, e.g., Crane 2003, Douglas 1996, Dretske 1995, and Tye 1995.

pain in one's arm is not an experience of some kind of mental object, the pain, but rather an experience of one's arm, or a region of one's arm – typically, though not always,[3] the strained, injured, or diseased region – by means of the pain. This does not, of course, mean that we cannot be conscious of our pains; on the contrary, we usually *are* conscious of our pains. But pains are first and foremost the means by which we are conscious of a given bodily region, which means that the object of a pain is whatever bodily region the subject of the pain attends to, i.e., whatever bodily region she says hurts. If the subject complains of having a stomach ache, then the object of her complaint, and hence the object of her pain, is her stomach, or a region of her stomach; if she complains of having an earache, then the object of her pain is her ear, or a region of her ear; if she has a painful cut on her finger, then object of her pain is the injured region of her finger; and so on.

Now it is not unusual for people to say "I am in pain" (*sans phrase*), but such an utterance, and others like it, should not be taken to indicate that pain is, or can be, non-intentional. Notice, first, that the statements "I have a pain" and "I am in pain" typically occur in contexts in which the speaker's interlocutor(s) is aware of the injured or diseased bodily region to which the speaker is referring when she says, "I have a pain" or "I am in pain." Moreover, if my interlocutor says "I'm in a lot of pain," and the context is such that I am unaware of the bodily region to which she is referring, then I would probably ask: "Where does it hurt?" And if she persists in saying, simply, "I'm in a lot of pain," then I would be puzzled and wonder what she means. The point here is that pains always have a felt location; that is, pains are always experienced as located in a particular region of the body. Sometimes the location is "global" or more generalized, i.e., covers a wide bodily area or many discrete bodily areas at once, in which case one might say "I am in pain *all over*," or "It aches *everywhere*," or "I am in pain." But this latter expression is always elliptical for "I am experiencing (feeling, etc.) pain in *x*," which can be further transcribed, if the intentional view of pain is correct, as "I am conscious, in a painful way, of *x*."

Tye has recently developed a theory of consciousness that we might recruit in support of this conception of the nature of pain states. Tye's theory is predicated on the idea that conscious states are representational states. Accordingly, he suggests that "pains are sensory representations of bodily damage or disorder. More fully,

---

3. It bears emphasizing that pains can be directed at bodily areas that are not damaged, as in the case of *referred* pain. For example, it is not uncommon to experience a momentary headache upon drinking a cold beverage too quickly. When this occurs, however, it is not one's head, presumably, that undergoes a disturbance. The pain, in such cases, has been referred to a bodily area that might be perfectly healthy.

they are mechanical responses to the relevant bodily changes in the same way that basic visual sensations are mechanical responses to proximate visual stimuli" (1995, p. 113). To illustrate, he presents an extensive, intuitively appealing taxonomy of various pains sensations and how they track or represent different sorts of bodily disturbances. For instance, stabbing pains represent sudden disturbances or injuries, usually over a specific bodily region; twinges of pain indicate mild, brief disturbances; racking pains represent disturbances involving the stretching of certain internal bodily parts, usually muscles; and pricking pains represent disturbances that cover a very tiny bodily area. In each of these cases, according to Tye, "the subject of the pain undergoes a sensory representation of a certain sort of bodily disturbance. The disturbances vary with the pain" (1995, p. 113).

Tye adopts this approach to pain primarily because he believes that, if correct, it would solve a number of traditional problems associated with pain sensations. In particular, he argues that his proposal satisfactorily accounts for the problem of felt location, the problem of explaining why, even though pains are typically experienced as features tokened in a particular region of the body, they are not actually located there. He also contends that the intentional approach solves the problem of the phantom limb, the problem of explaining how it is possible for people to sometimes experience pain in appendages that have been lost or amputated. According to Tye, we can experience pains in nonexistent limbs just as we can hallucinate pink elephants and believe in and think about things that do not exist, such as unicorns and golden mountains. Pains and other conscious states are representational states that have the capacity to represent or be about objects and events that do not necessarily exist.

The intentional approach to pain appears to explain a peculiar feature of our pain vocabulary as well, namely, that there is an implicit *intensionality* – with an "s" – in statements of pain location. Briefly, whereas intentionality – with a "t" – is a property of mental states,[4] intensionality – with an "s" – is a property of certain logical and linguistic contexts. To demonstrate the existence of this feature in statements of pain location, Tye advances a critique of an argument from Block purportedly showing that terms that apply to both pain and physical objects assume a systematic difference in meaning when applied to pain. Block (1983, p. 517) invites us to consider the following argument:

(P1) The pain is in my fingertip.
(P2) The fingertip is in my mouth.
(C1) The pain is in my mouth.

4. Though not *only* of mental states: see Martin and Pfeifer 1986.

This argument is obviously invalid, and its invalidity derives, according to Block, from the fact that "in," when applied to pain, differs in meaning from its ordinary spatial location sense. However, Tye provides a structurally analogous inference showing that Block's argument does not establish that "in" assumes a special meaning when applied to mental particulars. Here is the inference Tye (1995, p. 113) propounds:

(P1) I want to be in City Hall.
(P2) City Hall is in a ghetto.
(C1) I want to be in a ghetto.

This argument is also invalid, but the term "in" has the same meaning in both the premises and the conclusion. This indicates, according to Tye, that the term "in" does not systematically differ in meaning when applied to pain. Rather, the term "in," in both the first premise and the conclusion of the above argument, appears in an intensional context created by the "I want" clause. Tye suggests that this is precisely the way "in" operates in Block's inference. It is reasonable to expect, he concludes, that there is an implicit intensionality in all statements of pain location.[5]

Even if persuasive, the claim that statements of pain location create intensional contexts might seem puzzling. Clauses like "I want," "I believe," "I hope," etc. create paradigmatic intensional contexts because the contexts in which they appear fail two of the most common tests for extensionality; that is: (1) they do not license existential inference; and (2) they are referentially opaque – substitution of co-referential terms (or identicals) may produce a change in the statement's truth-value.[6] Statements of pain location, however, do not fail this latter test for extensionality, i.e., they do not exhibit the feature of referential opacity. For example, if "foot" is replaced with the co-referring noun phrase "lower extremity of the vertebrate leg," then the truth-value of the statement "I have a pain in my foot" does not change. Someone experiencing a pain in her foot *is* experiencing a pain in the lower extremity of her vertebrate leg.

On the other hand, although statements of pain location seem not to pass the substitutivity test for being intensional, they do create contexts that prevent

5. Carruthers (2000, pp. 118–21) agrees with Tye on this.

6. To be sure, it is has been claimed that contexts created by clauses like "I want," "I believe," "I hope" etc. do not fail this latter test for extensionality. Those who defend the view that the semantic content of a proper name, demonstrative, indexical, etc. is its referent are prepared to accept the view that co-referring expressions are intersubstitutable *salva veritate*, even in such contexts. I cannot explore this issue here.

existential inference. If, after a horrible accident, Jones awakens from a coma and says, "I have a pain in my foot," it does not follow that Jones has a foot. For suppose, unbeknownst to Jones, his foot was amputated while he was in a coma. In this case Jones is experiencing pain in a nonexistent limb. So the statement "I have a pain in my foot" does not, by itself, permit an existential inference with respect to a foot, just as the statement "I believe that there is a tooth fairy" does not, by itself, permit an existential inference with respect to a tooth fairy.[7] On this basis alone, we can conclude that statements of pain location create intensional contexts.

The fact that statements of pain location create intensional contexts should not be surprising. After all, statements of pain location, just like statements expressing beliefs, hopes, fears, etc. are, if anything, psychological statements. And, at the risk of belabouring the obvious, psychological statements are prime candidates for creating intensional contexts.

Moving from intensionality (with an "s") to intentionality (with a "t") might still occasion protest on the grounds that former is a property of certain logical and linguistic contexts, while the latter is a property of mental states. It might be objected, therefore, that moving from one to the other is illegitimate because such a move unwarrantedly presupposes an intimate connection between these two very different phenomena. Tye might be accused, in other words, of not heeding Searle's warning that "one of the most pervasive confusions in contemporary philosophy is the mistaken belief that there is some close connection, perhaps even an identity, between Intensionality-with-an-s and Intentionality-with-a-t" (1979, p. 85). As I see it, however, Tye's reasoning is as follows. There are certain peculiar intensional features of our pain vocabulary, and conceiving of pains as intent*i*onal *explains* these peculiar features more convincingly than Block's alternative. We might adopt the position that the term "in" assumes a systematic difference in meaning when applied to mental particulars, but such a position seems suspiciously ad hoc, at least in the absence of a plausible account of why the term assumes a special meaning only when applied to mental particulars. Thus, Tye does draw a connection between intensionality and intent*i*onality, but he does not identify them. The connection he draws is an explanatory one.

---

7. "I have a pain in my foot" and "I believe in the tooth fairy" are first-person psychological ascriptions, but third-person psychological ascriptions, such as "Jones has a pain in his foot" and "Jones believes in the tooth fairy," create intensional contexts as well, inasmuch as, like first-person psychological ascriptions, they do not licence existential inference. Strictly speaking, one cannot infer, simply on the basis of the statement "Jones has a pain in his foot," that Jones has a foot, just as one cannot infer, simply on the basis of the statement "Jones believes in the tooth fairy," that there is a tooth fairy.

Now, for reasons I will disclose in a later chapter (ch. 7), I am inclined to think that Tye's representational theory, and representationalism about consciousness in general, faces insurmountable difficulties. But Tye is able to account for the peculiarities of our pain vocabulary and solve the other puzzles regarding pain location because he treats pains as intentional. And the assumption that one can treat pains as intentional only if one treats them as representational has no obvious demand on our acceptance.[8] The point here is that we can move from intensionality to intentionality in the same way Tye does, and reap the explanatory benefits derived therefrom, without, thereby, incurring the difficulties afflicting representationalism. Thus, while I do not agree with Tye that pain is representational, it is plausible that there is an implicit intensionality in statements of pain location. And on the basis of the intensional reading of such statements, I would draw the same conclusion Tye draws, namely, that pain is intentional.

The view of pains as intentional can be readily extended to bodily sensations generally. For example, a tickle, like a pain, is a conscious state whose object is a bodily disturbance, a disturbance most commonly caused by the presence of an object lightly touching the skin, while an itch is a conscious state whose object is a more irritating and abrasive surface disturbance (see Tye 1995, pp. 116–19). In each case the intentional object is a bodily region that has undergone a surface disturbance of some sort.

Emotions, it seems, also have intentional content. Tye proposes that a sudden feeling of fear, dread, or anger is a conscious state whose object is constituted by certain physical changes in the body, such as rising blood pressure, an increasing heart rate, and rapid alterations in the immune system (Tye 1995, pp. 125–31). During emotionally tumultuous experiences we are conscious of the relevant bodily changes. Damasio's study of the neurological underpinnings of emotion offers some strong support for Tye's proposal. When emotions occur, Damasio claims,

> there is a change in your body state defined by several modifications in different body regions. If you meet an old friend. . .your heart may race, your skin may flush, the muscles in your face change around the mouth and eyes to design a

---

8. This is controversial. Some philosophers think that intentionality *is* representation (e.g., Fodor 1987, p. xi; Tye 1995, p. 95; Rey 1998, p. 435; Kim 1998, p. 21; Lycan 2001, p. 4); that is, they think that "a state has a thing as its intentional object only if it represents that thing" (Lycan 2001, p. 4). Here the claim is not that a state has a thing as its intentional object only if it is *about* that thing; rather, the claim is, crudely, that a state has a thing as its intentional object only if the state in question is an internal mental representation that *conveys information* about that thing. I offer a critical discussion of this position in Chapter Seven.

> happy expression, and muscles elsewhere will relax. . . . A number of peptide modulators are released by the brain into the bloodstream. The immune system also is modified rapidly.
>
> (1994, p. 135)

Tye, drawing on the work of Damasio, argues that background feelings, such as moods, have intentional content as well. Moods, like emotions, on his view, are mental states whose object is constituted by certain bodily changes, though, as he admits, it is sometimes difficult to identify what the pertinent bodily change is. He develops an argument stating, roughly, that the object of a mood is a *departure* from the physical states constituting our *functional balance*. When functional balance is present, we function in a harmonious way without feeling any particular mood. When moods descend on us, we are responding in a sensory way to a departure from the relevant range of physical states. We are sensing physical changes in our "body landscapes."[9] The intentional object of a felt mood, then, according to Tye, is not a well specified bodily region or process. Instead, moods are mental states whose objects are constituted by more general bodily changes, the effect of sensory input and output that is inconsonant with the physical states constituting functional harmony. As Damasio puts it: "As a whole, the set of alterations defines a profile of departures from a range of average states corresponding to functional balance, homeostasis, within which the organism's economy operates probably at its best, with lesser energy expenditure and simpler and faster adjustments" (1994, p. 135).

Despite its surface plausibility, Tye's account of the intentionality of moods and emotions is not without its difficulties. As a lead up to my remarks on these difficulties, I want to make a point of clarification. Tye neglects to adequately distinguish emotional perturbations and short-term affective disturbances – in short *feelings* – from longer-term dispositional states or moods *simpliciter*. Presumably, Tye is concerned with the former, not the latter, since the latter are not conscious states at all. In general, a dispositional state or a mood *simpliciter* is a "*proneness* to feel, during one's waking hours, joyful, or depressed, or cheerful, and so forth" (Bennett and Hacker 2003, p. 202, my emphasis). But it is not the proneness itself that is conscious, but rather the *feeling* of joyfulness, depression, cheerfulness, and so forth that issues from the proneness. Tye is also not concerned, presumably, with character traits, such as irascibility, malice, benevolence, friendliness, and the like, which, once again, are not feelings but rather dispositions to feel or act in a certain way. In short, the proper subject matter of Tye's intentional analysis of moods and emotions is (sometimes persistent) occurrent emotional and moody

9. Tye takes this phrase from Damasio 1994.

feelings, such as a feeling of happiness, a feeling of sorrow, a feeling of melancholy, a feeling of fearfulness, and so on.

If so, however, it is implausible that the intentional object of a mood or emotion is always a bodily change of some sort. The bodily change that accompanies a felt mood or emotion is, in many cases, more appropriately understood as the *effect* of that mood or emotion, not the object. If sleep deprivation (or ingestion of a mind-altering drug, etc.) is the source of the mood or emotion, then perhaps it is reasonable to maintain that the object of the mood or emotion is a bodily change of some sort. But if a person is feeling angry or depressed because he has lost his job, then the object of his anger or depression is the loss of his job. He may undergo certain physical changes as a result of his anger or depression, such as increased heart rate, rising blood pressure, etc. But these physical changes are caused by his anger or depression, not vice versa.

Moreover, felt moods, such as a feeling of depression, do not always have an *identifiable* object. One can feel depressed without knowing what one is feeling depressed about.[10] Like Tye, I would not conclude from this that felt moods are objectless, but neither would I conclude that they are *always* conscious states whose objects are constituted by certain of our physical states. We are usually able to identify the source of our moods (at least upon reflection), and the object of a felt mood is its source. If I am elated because an old friend is coming to visit, then the source of my elation, and hence the object of my mood, is my friend's impending visit. And if I cannot identify the source of my mood – if, for example, a feeling of depression descends on me for no discernible reason – then perhaps the source of the felt mood is certain of my physical states, in which case I am conscious, via my feeling of depression, of these physical states. However, the point is that it does not seem accurate to affirm that the object of a felt mood is always a bodily change of some sort.

So far the evidence seems to indicate that feelings and bodily sensations, conscious experiences typically regarded as non-intentional, can be reasonably construed as having intentional content. This view may not comport very well with out pre-theoretic intuitions – for example, felt moods and emotions, pains, and other bodily sensations do not initially seem to have an intentional object – but the seeming non-directedness of certain mental states is, I think, at least partly attributable to the lack of an identifiable intentional object at which these states

10. It might be appropriate to register a small disclaimer at this point. I do not pretend to have the resources to pronounce on what clinical psychologists refer to as "clinical depression." My concern here is with the ordinary concept of depression, and it is reasonable to suppose that one cannot extrapolate from the ordinary usage of this concept to more technical usages.

are directed. As indicated above, it is not always easy determine what the object of a mood (emotion, etc.) is. Sometimes we experience moods that have a definite object, and sometimes we experience moods that have, as Baier says, "all-inclusive and undifferentiated objects" (1990, p. 3).

At any rate, the view that all conscious states have intentionality has a number of attractive explanatory payoffs, which is reason enough to take the view seriously. As explained above, conceiving of pains as intentional nicely accounts for the peculiarities of our pain vocabulary. In addition, the intentional approach accounts for the problem of felt location. We experience pains as tokened in a particular region of the body when they are not actually located there because pains have intentional content; that is, pains are directed at whatever bodily region is strained, injured, diseased, or otherwise damaged. Pains, to be sure, can be directed at objects that do not exist, as in the case of the phantom limb (or be directed at bodily areas that are not damaged, as in the case of referred pain). But if pains have intentional content, then the problem of the phantom limb is just one aspect of the problem of explaining how one can be conscious of a nonexistent object, in this case an appendage. The problem of the phantom limb, then, just like the problem of felt location, is part of the larger problem of intentionality. This is not to suggest that we have reached a consensus with respect to the problem of intentionality, but we have a general idea what the problem is, and we can draw on the creative and ingenious efforts that have been put into understanding it (e.g., Searle 1979, 1983; Martin and Pfeifer 1986[11]). Moreover, if we can subsume the problem of felt location (the problem

11. An important point to be gathered from Martin and Pfeifer's article is that intentionality cannot be considered an exclusive property of mentality. Almost all of the familiar characterizations of intentionality hold for non-mental dispositional physical states as well. Intentional mental states are typically characterized in terms of that (which may or may not be present or exist) to which they are directed. But, as Martin and Pfeifer point out, a non-mental physical dispositional state needs, just as much as does a mental state, to be characterized in terms of that to which it is directed. To take one of Martin and Pfeifer's examples, consider the obvious parallel between the mental state ascribed to the subject in "Jake fears lions," and the physical disposition ascribed to the subject in "X is soluble in aqua regia" (Martin and Pfeifer 1986, p. 533). The mental state ascribed to the subject in "Jake fears lions" requires characterization in terms of that to which the mental state is directed. Similarly, the physical disposition ascribed to the subject in "X is soluble in aqua regia" requires characterization in terms of that to which the disposition is directed. Observe, as well, that the physical disposition in "X is soluble in aqua regia" is directive and *selective*. It is a disposition *for* one thing rather than another. Consequently, "Just as a belief needs *content*, namely, what would follow the 'that' in 'belief that. . .' so a disposition needs a *what-for*, namely, what would follow the 'for' in 'disposition for. . .', e.g., 'disposition *for* dissolving in H_2O' (and *not aqua regia* or xyz)" (Martin 1997, p. 211).

of the phantom limb, etc.) under the problem of intentionality, then we are being productive with regard to the goal of simplification and unification. Considerations of explanatory simplicity, then, seem to support the view that all conscious states have intentionality.

I do not take the arguments just adduced to be definitive. Such an object-oriented account might strike some as simply false of a feeling of dizziness (see Woodruff Smith 1989, p. 81), of our occasional states of reverie, and of the conscious states of the infant in the womb. In addition, the very general model I have presented might be thought to express intuitions about quintessentially empirical questions, questions to which scientific answers are now only forming. However, diverging at greater length on the topic of whether all conscious states are intentional would take us too far afield. Suffice it to say, the view is eminently defensible, profiting from a convergence of linguistic, phenomenological, and neurophysiological evidence.

4.3 The Transitivity Principle

I will now defend, and offer the promised direct arguments for, the thesis that a conscious state is always a state whose subject is, in some way, aware (or conscious) of being in it.[12] Following Rosenthal (e.g., 2000b, 2005), I will refer to this thesis as the Transitivity Principle. It is worth noting that some authors who are sympathetic to the thesis deem it unnecessary to defend it at all. Lycan (2001), for example, employs it as a stipulative definition in his argument for the claim that some higher-order representation theory of state consciousness is correct. This attitude might strike some as overly cavalier, but I am inclined to think that it is entirely justified, especially in light of the fact that there is a venerable philosophical tradition, dating back to at least Aristotle, that takes the thesis more or less for granted. Still, some notable philosophers have recently criticized it, so I feel it best to defend it.

One of those philosophers is Dretske, so I will begin by fending off his line of resistance to the thesis.[13] Dretske has argued that *any* account of state consciousness

---

12. I will take the following two theses to be equivalent: (1) a conscious state is a state whose subject is conscious (or aware) of *being in* it; and (2) a conscious state is a state whose subject is conscious *of* it. There are subtle differences between these theses, but for present purposes they can be ignored.

13. Another of the Transitivity Principle's detractors is Siewert, who maintains that we are led to the view that conscious states necessarily involve self-awareness "by an insufficiently critical reliance on certain ways of speaking" (1998, p. 197). Among others who believe that state

that takes the consciousness of a state to lie in the awareness *of* that state rests on a confusion about what it is for a state to be conscious. He states his view as follows:

> This conclusion [that a subject's awareness of her mental state makes that state conscious] is most peculiar. It is not to be accepted lightly. It is, in fact, not to be accepted *at all*. There are, to be sure, states in (or of) us without which we would not be conscious of trees and pianos. We call these states experiences. Since these experiences make us conscious of things (pianos, trees, French horns) the states themselves can be described as conscious. But we must be careful not to conclude from this that because the states are conscious, we must, perforce, be conscious of them. That doesn't follow.[14]
>
> (1995, p. 100)

Dretske's argument for this claim can be sketched as follows. Suppose Clyde looks at a photograph of twenty red roses at *t* and then looks at the same photograph at *t*+, failing to notice at *t* that, between *t* and *t*+, one of the roses was airbrushed out of the photograph. On Dretske's view, since Clyde saw a photograph of twenty roses at *t* and a photograph of nineteen roses at *t*+, Clyde's conscious visual experience at *t* differs from his conscious visual experience at *t*+. Nevertheless, Clyde is not *aware* that his experiences differ.[15] This indicates, Dretske claims,

> . . .that there can be conscious differences in a person's experience of the world – and, in this sense, conscious features of his experience – of which that person is not conscious. If this is true, then it cannot be a person's awareness of a mental state that makes that state conscious. . . . It follows, therefore, that what *makes* a mental state conscious cannot be our consciousness of it. If we have conscious experiences, beliefs, desires, and fears, it cannot be our introspective awareness of them that makes them conscious.
>
> (1993, pp. 278–9)

This argument, however, is unsound: we may dispute the assumption that Clyde's conscious visual experience at *t* differs from his conscious visual experience

---

consciousness is possible without self-awareness are Bennett (1966), S. White (1987), and Searle (1992, 1999). If the considerations adduced below are correct, however, the prospects for a theory of consciousness that does not take state consciousness to necessarily involve self-awareness are unpromising at best.

14. A minor terminological qualification should be made here. Dretske equates states with experiences in this passage, which may seem odd, since states are things we can be said to be *in*, whereas experiences are things we *have*. Dretske's point, though, is that if one is in a conscious perceptual state, then one is having a perceptual experience; if one is in a conscious emotional state, then one is having an emotional experience; and so on.

15. This is a distilled version of an example found in Dretske 1993, pp. 272–8.

at *t+*. The *object* (or *content*) of Clyde's visual experience at *t+* differs, of course, from the object of Clyde's visual experience at *t*, but it does not follow that Clyde's *visual experience* at *t+* differs from Clyde's visual experience at *t*. Dretske claims that two experiences must be held to differ because they "are not only experiences of [a] qualitatively different object [a photograph of twenty roses at *t* and a photograph of nineteen roses at *t+*], they are experiences of the qualitative differences. The respects in which [the photograph] differ[s] are not only visible, they are (by hypothesis) seen" (1993, p. 274). But are they? All can agree that Clyde would deny that he had seen any difference in the photograph; and if asked how many roses he saw (at both *t* and *t+*) Clyde would say "Many" or "Several," and this is the natural, indeed only, response to be given here (assuming, of course, that Clyde did not count the number roses in the photograph). So why should we suppose that Clyde has misreported his own visual experience? To be sure, there are a determinate number of roses in the photograph, but it does not follow that Clyde *consciously experienced* a determinate number of roses. Brentano has claimed that even if a strong physical stimulus produces a visual experience (auditory experience, etc.), a weaker stimulus need not necessarily result in a visual experience of lesser intensity (1874, p. 117). He makes this claim in the course of attempting to refute various attempts to establish the existence of unconscious mental states. His opponents here are, among others, Leibniz and Kant, who argued that we must be conscious of (unconsciously aware of?) minute visibilia, audibilia, etc. – items Kant called "obscure representations" (1965, B414-15a) – that we cannot report. According to Leibniz and Kant, if you look outside your window at a snowfall, you see, if only unconsciously, every snowflake that falls within your visual field, since your visual experience can only be composed of the visual sensations produced by the individual snowflakes. According to Brentano, however, an obvious fallacy vitiates this argument. "The effect of a sum of forces," he claims, "is not only quantitatively, but very often also qualitatively different from the effect of its individual components. If the temperature drops to just above zero degrees it does not change water either partially or to a smaller degree into ice" (ibid.). Along similar lines, it might plausibly be claimed that even if a strong physical stimulus, say a photograph of twenty roses, produces a visual experience, a "weaker" stimulus, say the same photograph with an item missing, need not necessarily result in a different visual experience.

Now Dretske, as I briefly mentioned in Chapter One, holds that it is possible to be conscious *of* a thing, *x*, without thereby being conscious *that* it is an *x*. He calls the former "thing awareness" and the latter "fact awareness." It might be argued, on the basis of this distinction, that when Clyde looks at the photograph at *t+*, Clyde is thing aware of the difference in the photograph, but not fact aware; that is, he is in a conscious visual state in virtue of which he is aware *of* the difference (despite

his protestations to the contrary), but he is not in a conscious visual state in virtue of which he is aware *that* there is a difference (or of the fact that there is a difference). Thus, Clyde's visual experience at *t*+ differs from Clyde's visual experience at *t*, since at *t*+ Clyde is not in visual state in virtue of which he is aware of the difference in the photograph, whereas as *t*+ he is.

This argument is unconvincing. I see no warrant for the claim that at *t*+ Clyde is in a conscious visual state in virtue of which he is aware of the difference in the photograph. Notice that the sort of case at issue here is not like the sort of case in which, in A. R. White's words,

> we significantly say not that someone 'did' see so and so without realising it, but that he 'must' have seen it. . . . Here we are claiming that he saw something somewhere without realising that he had seen anything there. The reason we cannot say that he 'did' see it is that what we say he 'must' have seen is not what he did admittedly and knowingly see under another description and what, therefore, it logically follows he did see, but something which we have various other reasons, good or bad, for arguing that he 'must' have seen. . . . For instance, if something fairly prominent is right in front of one, say a big bush beside a shed, but I do not remember, even immediately afterwards, having seen it, it might be true in different instances to say either that I must have seen it though I did not realise I had or that, strange though it may seem, I did not in fact see it.
> (1964, p. 54)

There is no evidence one could produce to support the claim that Clyde *must* have been in a mental state in virtue of which he consciously experienced the difference in the photograph. To be sure, Clyde was in a mental state in virtue of which he consciously experienced, in this case visually, *the photograph* at *t*+, but we consciously experience a lot of things at different times without being in mental states in virtue of which we consciously experience all the minute differences these items undergo.[16]

Speaking at a more general level, the temptation to think that Clyde's visual experience at *t* differs from his visual experience at *t*+ derives, at least in part I think, from tacit acceptance of the canard, outlined in Chapter One, that the way a conscious experience is in itself might transcend the way it is for the experiencing subject, i.e., the way phenomenology reveals it to be. To claim that Clyde saw the difference in the photograph is to assume that Clyde's visual experience has, in

16. Cf. Merleau-Ponty: "A really homogeneous area, offering *nothing to perception* cannot be given to *any perception*" (1962, p. 4). What Merleau-Ponty means here, according to Dwyer, is that "a putative item of perception, insufficiently articulatable to engage the discriminatory power of perception, is in fact imperceptible" (Dwyer 1990, p. 11).

itself, a determinacy that it does not have for Clyde. But Clyde, on each occasion, saw a photograph of a *bunch* of roses, and this crucial feature of Clyde's visual experience, i.e., its indeterminacy,[17] cannot be reckoned to be a *seeming* experiential property – there is no seeming about it – that a non-experiential property might transcend. To suppose otherwise betokens a failure to appreciate that it makes no sense to say that the way a visual experience (or any other experience) is for a subject might be another way in itself.

It will be useful to frame this issue in terms of the distinction, previously considered (see Chapter One), between *personal* and *sub-personal* levels of content. The kind of content that is pertinent here, i.e., the content of Clyde's *conscious visual experience*, is not sub-personal content (typified by, for example, the "implicit representations" of cognitive science and psychology), but rather personal-level (or conscious, or experiential) content: it "specifies the way the world appears or seems *to the subject*" (Byrne 2001, p. 201).[18] Consider the famous Müller-Lyer illusion, where a subject is presented with two lines of equal length, one with arrows directed outward and the other with arrows directed inward. Even though the lines are the same length, the personal-level content of the subject's visual experience is of two lines of unequal length, since *to the subject* the line with the arrows directed outward appears longer than the line with the arrows directed inward. Applying the distinction between personal and sub-personal levels of content to the case at hand, we can say that even though the sub-personal content of Clyde's visual experience at *t* may differ from the sub-personal content of his visual experience at *t+*, it does not follow that the personal-level content differs: the way the world appears *to Clyde* is the same at both *t* and *t+*. Otherwise put, the sub-personal content of Clyde's visual experience at *t* is (at least in part) a photograph of twenty roses, whereas the sub-personal content of his visual experience at *t+* is (at least in part) a photograph of nineteen roses. But the personal-level content of

17. Cf. A. R. White: "the mere facts that you saw X and realised that you saw it and that X is composed of visible parts are not sufficient reason for saying that you must, with or without realising it, have seen these parts. For to say you saw X may be to identify what you saw and there is no straightforward rule about how much of X you must see in order to be said to see X" (1964, p. 55).

18. There are a number of complicated questions regarding this notion of content – e.g., whether it is best treated along Russellian, Fregean, or possible worlds lines (for discussion of these treatments of content see, respectively, Peacocke 1992, Woodruff Smith 1989, and Stalnaker 1999); whether it ought to be individuated widely in a way that is sensitive to social and environmental factors; and so on. My argument in this section does not require that I take up a position on any of these controversies.

his experience is indeterminate: what is presented to Clyde at both *t* and *t+* is a photograph of a *bunch* of roses.

Allow me to pause momentarily to take stock. Dretske has argued that since there can be conscious differences in a subject's experience of the world of which the subject is unaware, it cannot be the case that what makes a mental state conscious is the subject's awareness of it. However, Dretske's argument relies crucially on a false assumption, namely, that minute differences in consciously perceived visibilia, audibilia, etc. necessarily result in corresponding changes in one's conscious perceptual experiences of those visibilia, audibilia, etc. And since this assumption is false, Dretske's argument is unsound.

4.4 Two positive arguments for the Transitivity Principle

Having motivated dissatisfaction with Dretske's argument against the Transitivity Principle, I want now to advance two positive arguments for the thesis. I will call these arguments the Symmetry Argument and the Argument from Spontaneous Reportability.[19] As far as I know, both arguments are novel,[20] at least as arguments purporting to demonstrate the truth of the Transitivity Principle, a thesis for which there is, in general, a dearth of direct arguments. To be sure, a number of philosophers, including Brentano (1874), Sartre (1956, 1957), and, more recently, Rosenthal (e.g., 1986, 1990, 2000b), espouse one version or another of the thesis, but generally these philosophers either assume its truth on phenomenological grounds or provide indirect arguments for it by arguing for the falsity of its negation (e.g., Rosenthal 2000b, §III).

Of the small number of direct arguments for the thesis, Sartre's is probably the most well known. Sartre puts forward a *reductio*, arguing that since all consciousness is consciousness *of* something, if one's consciousness of something were not a consciousness of itself, then it would be a consciousness ignorant of itself, an unconscious consciousness, which, according to Sartre, is absurd (1956, p. lii).[21]

---

19. I should note that the thesis is *not* that all conscious states necessarily involve an explicit or reflective awareness. For more on this point, see below.

20. Though Sartre (1956) and Rosenthal (1990, 1993d, 2002c) have propounded arguments that bear some resemblance to the argument from spontaneous reportability.

21. Cf. the following argument from Gurwitsch, which is a virtual restatement of Sartre's argument: "Let us suppose an object to be given, i.e., an act to be experienced through which the object is presented. . ., however, without our being aware of the act. This would amount to the

I find this argument persuasive; in fact, I propounded a version of it in Chapter Three, when I argued that conscious states that are entirely first-person unknowable cannot be said to be conscious states. Another, less well-known (at least in the west), argument for the Transitivity Principle is the memory argument from the Indian Buddhist philosopher Dinnaga (see Ganeri 1999). This argument, to oversimplify enormously, states that since no past event can be remembered unless it has been previously experienced, it is reasonable to conclude that no past experience can be remembered unless it has been previously experienced. But to experience an experience is to be aware of it; thus, we are aware of our experiences. In my judgement, this argument is as compelling as Sartre's *reductio*. Considered in tandem, then, Sartre's *reductio*, Dinnagā's memory argument, the Symmetry Argument, and the Argument from Spontaneous Reportability should provide powerful support for the Transitivity Principle.

4.4.1 The Symmetry Argument

The Symmetry Argument may be sketched as follows. As I have said, the expression "There is something it is like to V" (and its cognates), where V is some psychological verb or verb phrase, is always shorthand for "There is something it is like for S to V," where S is some creature. Conscious experiences, in other words, are always like something *for the subject* or *to the subject* of the experience. Observe that it is not inappropriate to say that ordinary objects and events are like something to us as well, in the sense that they look, taste, feel, etc. like something. Locutions of the form "*x* looks (feels, smells, etc.) so-and-so to S" – where "looks," "feels," "smells," etc. are, to borrow Chisholm's (1957) terminology, *appear* words – are used to express this idea. If *x* looks so-and-so to S, then *x* is like something to S. Chisholm (1957, ch. 4) points out that locutions of the form "*x* looks (feels, smells, etc.) so-and-so to S" express the *non-comparative* use of appear words. He contrasts locutions of this form with locutions of the form "*x* looks like a so-and-so to S" and "*x* looks to be so-and-so to S." Locutions of the form "*x* looks like a so-and-so to S" express the *comparative* use of appear words, while locutions of the form "*x* looks to be so-and-so to S" express the *epistemological* use of appear words. Locutions of the form "*x* looks so-and-so to S," in contrast, do not make a comparison between *x* and some other thing, nor do they make an epistemological claim to the effect that S is inclined to believe or has evidence to believe that *x* is so-and-so. The locution "*x* looks so-and-so to S" simply

---

assumption that we are confronted with an object and experience the corresponding act without having any knowledge at all of our experience; we would have consciousness of the object without being conscious of this fact. It is sufficient to contrive this assumption in order to see at once its absurdity" (1950, p. 5).

says that S perceives *x* as having such-and-such a property, i.e., that such-and-such a property constitutes *how x* is to S or what *x* is *like* to S. If *x* is a tree, for example, then *x* might look green, feel rough, smell mildewy, etc., and these properties constitute what *x* is like to S. The greenness, the roughness, the mildewy odour, etc. are, to be sure, properties of *x*, not properties of S's perceptual experience of *x*,[22] but these properties enter into a description of how *x* looks, feels, smells, etc. to S. There is, then, a symmetry between facts about what ordinary physical objects and events are like and facts about what conscious experiences are like (to us).[23]

Now I take it as self-evident that an ordinary object (event, etc.) cannot be like something to a subject, i.e., cannot look, feel, smell, etc. like something to a subject, unless the subject is aware, in one way or another, of that object. For example, in order for a tree to look like or seem like something to me, I have to perceive the tree, imagine the tree, hallucinate the tree, or otherwise be aware of the tree. So I take the following principle as obvious:

(1) Physical object *x* is like something to S at *t* ↔ S is aware[24] of *x* at *t*.

The same, I suggest, is true of conscious experiences: in order for a conscious experience to be like something to the subject of the experience, the subject must, in some way and to some degree, be aware of it. That is:

(2) Conscious experience *y* is like something to S at *t* ↔ S is aware of *y* at *t*.[25]

---

22. It can be true that *x* looks so-and-so to S even if there is no *x*, in which case it might be more appropriate to say that the properties ascribed to *x* are, in some sense, properties of S's experience of *x*.

23. This does not mean that conscious experiences instantiate extra-mental properties, or that they instantiate greenish (mildewy, etc.) "feels" or "sensations." The symmetry here issues from the fact that extra-mental objects (events, etc.), just like conscious experiences, are like something to us.

24. This use of "aware" deserves comment. It might be argued that, strictly speaking, the biconditional expressed here is false, since S can be aware of *x* even though it is not the case that *x* is like something to S. For example, a sleepwalker is aware of her surroundings, but they are not like anything to her. As I have already noted, however, for the purposes of this book I regard "aware" and "conscious" as synonyms, and so any occurrence of "aware" can be replaced with "consciously aware." And if "aware" is replaced with "consciously aware," then the biconditional is obviously true.

25. I should remind the reader that insofar as it is not possible to be conscious that F without, thereby, being conscious *of* F, I am subsuming being conscious "*that* something is so" under the rubric "conscious *of* something."

And since, as I argued at some length in Chapter Three, all conscious experiences are, indeed, like something to the subject of the experience, it follows that the subject is, in some way and to some degree, aware of the experience. So just as a physical object cannot be like something to a creature unless the creature is aware of that object, so a conscious experience cannot be like something to a creature unless the creature is aware of that experience. It follows that a conscious experience is always an experience of which the subject is, in some appropriate way, aware.[26]

Two points bear emphasizing. First, if the Symmetry Argument is sound, then so-called "first-order" accounts of state consciousness (e.g., Dretske 1995, Tye 1995) – i.e., accounts of state consciousness that do not construe a mental state's being conscious in terms of the subject's awareness of it, but rather in terms of the mental state itself making the subject conscious of some object – are inadequate. However, the Symmetry Argument does not specify what kind of self-awareness is involved in conscious states. It says (merely) that a conscious experience is like something to a subject just in case the subject is, in some suitable way, aware of that experience. But this conclusion leaves it open whether the subject's awareness of his or her experience is a reflective or non-reflective awareness, or even an unconscious awareness. As formulated, then, the Symmetry Argument cannot be pressed into service to support any specific theory of consciousness that takes state consciousness to entail self-awareness. Rather, its purpose is to answer the preliminary question of whether or not state consciousness does in fact entail self-awareness.

Second, the symmetry between (1) and (2) issues from the fact that the object *x*, just like the conscious experience *y*, is *like something* to S, not that *x*, like *y*, is an object. The Symmetry Argument does not assume that only objects can be like something to a subject. I underscore this point to forestall the impression that the argument is designed to bolster the view that the kind of self-awareness involved in conscious states is of the "subject-object" variety, a view to which a number of phenomenologists, including Husserl and Sartre, have expressed opposition. (For detailed discussion of this point, see Chapter Six, §6.2.2.) The Husserlian/Sartrean

---

26. The pattern of reasoning here can be represented as follows:

(1) *X* is like something to S if and only if S is aware of *x*.
(2) All conscious states are such that there is something it is like to be in them.
(3) Therefore, all conscious states are states whose subject is, in some appropriate way, aware of being in them.

If this argument is sound, then the following two theses, held conjointly by many, are conjointly inconsistent: (i) all conscious states have phenomenal character, and (ii) there are conscious states whose subjects are entirely unaware of being in them.

view, however, that the self-awareness involved in conscious states is not to be construed in terms of a subject-object relation is entirely consistent with the Symmetry Argument, since, as I remarked, the argument does not assume that only objects can be like something to a subject.

4.4.2 Objections and replies

I will now confront a pair of objections that might be raised against the Symmetry Argument.

The first objection is that I have proceeded to hastily, that we cannot infer from the fact that a conscious experience *c* is like something to S that S is aware of *c*. One can be the subject of an experience, and, in many cases, one can reflect on and form various beliefs about one's experience. But only in these latter cases of introspection is there a symmetry between one's awareness s of one's experience and one's awareness of an external object. Thus, all we can infer from the fact that *c* is like something to S is that S *had c*, not that S was aware of *c*.

This objection is not to the point. We can assume that *c* is like something to S if and only if S has *c*. But this is platitudinous. The fact that having *c* is a necessary and sufficient condition for *c* to be like something to S fails to carry any pregnant information about the experiential "mechanism" responsible for the fact that *c* is like something to S, even if S does not reflect on *c*. In other words, what is it about having *c* that explains, or partly explains, the fact that *c* is like something to S even if S does not reflect on *c*? The Symmetry Argument provides a non-trivial response to this question. In having *c*, S is, in some sense, aware of *c*, and S's being aware of *c* partly explains the fact that *c* is like something to S, in much the same way that S's being aware of some extra-mental object *x* (event *x*, etc.) partly explains the fact that *x* is like something to S.

A simpler, yet no less effective, reply to this objection is to say that it strains credibility to suggest that in order for *any* extra-mental object or event to be like something to a subject, the subject must be aware of that object or event, whereas in order for a conscious experience to be like something to a subject, the subject need not be aware of it at all. It is the opponent of the Symmetry Argument who has some explaining to do regarding how there could be such an *asymmetry* between facts about what extra-mental objects are like and facts about what conscious experiences are like.[27]

---

27. As I pointed out in Chapter Two, Hacker has evinced scepticism about the very idea that we can define conscious experience in terms of what it is like to have an experience. But Hacker makes an observation that actually *supports* the symmetry argument. According to Hacker,

Another possible objection is that the argument trades on an equivocation. There is a crucial difference, it might be argued, between the case in which *x* is an external, non-mental object, and the case in which *x* is a conscious experience. Where *x* is an external, non-mental object, e.g., a tree, we need to specify a way in which S is conscious of *x* before the question of what the tree is like for S is even intelligible. The difference is such that the language of "what-it-is-like" has a different logic in the two cases. The tree, in itself and apart from the way in which S is aware of it, is not like anything for S. But this is not the case when *x* is a mental state, e.g., when *x* is S's visual perception of the tree. Thus, it is futile to attempt to draw any conclusions regarding the question of whether state consciousness entails self-awareness by appeal to our awareness of extra-mental objects.

This objection is likewise not to the point. When we ask what is required for something (whether physical object or mental state) to be like something to a subject, the following general response seems apt: "*X* is like something for S at *t* if and only if S aware of *x* at *t*." This has intuitive resonance where *x* is a physical object of some sort. And there is no reason to suppose that the logic of the case where *x* is an extra-mental object differs significantly from the logic of the case in which *x* is a mental state. To be sure, the Symmetry Argument does not, as stated above, say anything about what *kind* of self-awareness is involved in conscious states, but then it does not purport to. It says (merely) that a mental state is conscious if and only if the subject is, in some *appropriate* way, aware of being in it, i.e., that an appropriate awareness of a mental state M is a necessary and sufficient condition for M's being conscious. What kind of awareness is appropriate is a matter to be investigated.

4.4.3 The argument from spontaneous reportability

My second argument for the Transitivity Principle brings into play the disposition to spontaneously report on our own conscious experiences. It is not seriously contestable that in order to spontaneously report on an external object (event, etc.), one must, in one way or another, be (or have been) aware of it. And the same,

---

"Every experience is a *possible* subject of positive and negative attitudinal predicates, e.g., predicates of pleasure, interest, attraction" (2002, p. 165). And further along he says: "We can. . .ask with respect to an experience what it was like to undergo it, and the answer, if there is one, will specify whether one found it enjoyable or unpleasant, interesting or boring, frightening or exciting, etc." (2002, p. 174). These claims support the symmetry argument because just as one cannot find some worldly object, event, or state of affairs agreeable, disagreeable, delightful, etc. unless one is (or was), at least to some degree, aware of that object, event, or state of affairs, so one cannot find an experience agreeable, disagreeable, delightful, etc. unless one is (or was), at least to some degree, aware of that experience.

I surmise, is true of conscious experiences: in order to spontaneously report on a conscious experience, one must, in one way of another, be (or have been) aware of it. Thus, if I can spontaneously report on my conscious experience, then I am (or have been), aware of that experience. And since we *can* spontaneously report on our conscious experiences – and not just on the objects of those experiences (e.g., I can spontaneously report not only on the *rose* at which I am looking, but also on that fact that I am *looking* at the rose) – it is reasonable to conclude that we are (or have been) aware of them. Ordinarily, of course, we are not reflectively or introspectively aware of our experiences, but such experiences are nevertheless spontaneously reportable. If I smell a rose, for example, I may not reflect on my occurrent olfactory experience, but, if prompted, I could spontaneously report that I was experiencing the smell of a rose; and this seems possible only on the assumption that I was aware of my olfactory experience prior to being prompted to report on it. This is the point Sartre is getting at in the following memorable example:

> If I count the cigarettes which are in that case, I have the impression of disclosing an objective property of this collection of cigarettes: *they are a dozen*. This property appears to my consciousness as a property existing in the world. . . . If anyone questioned me, indeed, if anyone should ask, "What are you doing there?" I should reply at once, "I am counting." This reply aims not only at the instantaneous consciousness which I can achieve by reflection but at those fleeting consciousnesses which have passed without being reflected-on.
>
> (1956, p. 1iii)

Sartre's point is that even if I do not reflect on my act of counting, I can, if prompted, spontaneously report that I have been counting. But then this seems possible only if I was, in some way, aware of my act of counting prior to being prompted to report on it.

There is a possible difficulty with this proposal. So far I have focused my examples on perceptual and sensory states and conscious activities, but what about beliefs? Many beliefs, such as my belief that the Battle of Antietam was fought on September 17, 1862, are spontaneously reportable even though there is, it seems, no awareness of them whatever. But then it seems to follow that it is not the case that the spontaneous reportability of a mental state is explicable only on the assumption that the subject is, or was, aware of that mental state.

In response to this objection, let me speak, first, to the assumption that beliefs are mental states. Collins (1987), Malcolm (1991), Hacker (e.g., 1998), and Bennett and Hacker (2003, §10.2.1) have extracted arguments from the work of Wittgenstein that forcefully suggest that beliefs are not mental states at all, at least if "mental state" is understood in its ordinary sense. Mental states are things one can be said to be *in*; for example, we can say of someone that she is in a state

of reverie or a state of concentration. But we cannot intelligibly say of someone that she is in a state of believing, any more than we can say of someone that she is in a state of knowing. Moreover, mental states have what Wittgenstein called "genuine duration" (1967, §§71-85, 472). Beliefs, like mental states, have a beginning and (sometimes) an end, but beliefs lack genuine duration: one's belief that grass is green cannot be interrupted. "To believe something intermittently would be to believe it and then reject it, and subsequently to come to believe it again" (Hacker 1998, p. 83). What is more, mental states have characteristic behavioural manifestations and facial expressions. Unhappy people have unhappy demeanours, anxious people fidget, etc. Believing that *p*, on the other hand, does not have a characteristic behavioural manifestation or facial expression. Believing that *p* might give me reason to be unhappy or anxious, but "one cannot read off the set of my beliefs from my face and demeanour" (ibid., p. 84). Perhaps most damaging to the near-unanimous view that beliefs are mental states is the fact that, unlike mental states, we have an uncountable number of beliefs at any given time, and there are no criteria for "countably individuating beliefs" (ibid., p. 83). If I believe that the Earth is 150,000,000 kilometers from the sun, then I also believe that the Earth is at least 149,999,999 kilometres from the sun. But then I also believe that the Earth is at least 149,999,998 kilometres from the sun, at least 149,999,997 kilometers from the sun, at least 149,999,996 kilometres from the sun, and so on. Is each of these beliefs a mental state? If so, then I am now in over 150,000,000 mental states, which is absurd. I can, to be sure, be in two or more mental states simultaneously, e.g., I can be irritable and melancholy at the same time. But it makes no sense to say that, at this moment, I am in an indefinite number of mental states.[28]

If the preceding remarks are correct, then it follows that to report on one's belief that *p* is not to report on a mental state, just as to report on one's knowledge that *p* is not to report on a mental state.[29,30] But then this is hardly surprising. "I believe that *p*" is typically a hedged assertion of *p*,[31] and just as asserting that *p* is

28. Incidentally, Hacker (1998, pp. 84–9) gives equally compelling reasons for supposing that beliefs are not dispositional states.

29. Cf. Collins: "belief is much more like knowledge than philosophers have supposed" (1987, p. 33). Knowledge, to be sure, is factive whereas belief is not, but "belief exhibits the same kind of irreducible connectedness to the world outside the believer as knowledge" (ibid.).

30. Williamson (2000) has recently argued that knowing is a mental state, but his use of the term "mental state" is partly technical – i.e., stipulative; and since I am using "mental state" in its ordinary sense, my claim here does not implicate his view.

31. See Collins 1987, ch. 2.

not to assert anything mental, so asserting that *p*, along with a proviso that *p* might be false, is not to assert anything mental. If am asked about my belief that *p*, then I will give reasons for believing *that p*; that is, I will give reasons for believing that the content of my belief – i.e., whatever is expressed by the sentence "*p*" – is true. But to give reasons for believing that *p*, is to give reasons for why one deems *p* to be true; it is not to give reasons for why one is sometimes in a mental state involving *p*. Similarly, one can reflect on one's belief that *p*; for example, I can reflect on my belief that Kafka's "K" is a cipher. But to reflect on my belief that "K" is a cipher is to reflect on why I believe (or whether I should believe, etc.) that "K" is a cipher, and this involves weighing the evidence for and against the proposition *that "K" is a cipher*. It does not involve weighing the evidence for and against a mental state that I did not possess prior to forming the opinion that "K" is a cipher.[32]

Of course, these reflections do not tell against the objection, for one may concede that beliefs are not mental states and still maintain that the Argument from Spontaneously Reportability is unsound. Recall that the Argument from Spontaneously Reportability relies on the assumption that in order to spontaneously report on an extra-mental object, one must, in one way or another, be (or have been) aware of it; and the objection can be reformulated as an attack on this assumption. The objection, so reformulated, is not that it is possible to spontaneously report on a *mental state* of which is not, nor ever has been, aware, but rather that it is possible to spontaneously report on *something* – e.g., the content of a belief – of which one is not, nor ever has been, aware.

It is by no means obvious, however, that it is possible to spontaneously report on the content of a belief without being, or ever having been, aware of the content of that belief. If I report on the content of my belief, say my belief that *p*, then I report that *p*, along with a rider that *p* might be false. But to do this I have to think that (or about) *p*, however momentarily, and this means that I am aware, however momentarily, that (or of) *p* (just as to report on my knowledge that *p*, I have to think, however momentarily, that *p*, which means that I am aware, however momentarily, that *p*). I could not report on my belief that *p* (or on my knowledge that *p*) without thinking, at least briefly, that *p*. Consider, for example, my report that I believe that "K" is a cipher. If this is a genuine report, and not just a random utterance, then I could not have made the report without thinking that "K" is a cipher. I do not, as I said, think about a mental state and subsequently report on a mental state when I report that I believe that "K" is a cipher. Rather, I have the thought that "K" is a

32. For more on this point, see Bennett and Hacker 2003, §10.2.1.

cipher, and so become conscious that, and subsequently report that, *"K" is a cipher* (which is not a mental state). So my successful report that I believe that "K" is a cipher relies on my being aware, however momentarily, that "K" is a cipher.

Now it has been claimed that we should "resist the idea that having thoughts about objects we take to be distant in place or time, such as Saturn or Caesar, makes one conscious of those objects. . . . [Additionally] doubting and wondering something about an object do not make one conscious of the object" (Rosenthal 2004b, p. 159). This seems right, but it does not undercut my argument. If I think that Brutus killed Caesar, I do not become aware of either *Brutus* or *Caesar*, but, in thinking that Brutus killed Caesar, I occupy a conscious state that has content; and the content of my conscious state is *that Brutus killed Caesar*. Thus, it is reasonable to hold that, in thinking that Brutus killed Caesar, I am aware, however momentarily, of a proposition: the proposition *that Brutus killed Caesar*.

Let me stave off a possible misconception by pointing out that the Argument from Spontaneous Reportability is not a thinly veiled endorsement of the view that verbalization or reportability is a *criterion* for consciousness. Dennett seems sympathetic to such a view, contending that "a hallmark of states of human consciousness is that they can be reported (barring aphasia, paralysis, or being bound and gagged, for instance)" (1991, p. 307). But such a view seems to rule out *a priori* the existence of conscious states in infants and animals, which is implausible. The Argument from Spontaneous Reportability, then, is best construed as an argument that *reinforces*, rather than underpins, the case for the Transitivity Principle. The idea is that a pervasive feature of conscious states, at least of those of linguistic creatures, is that they are spontaneously reportable, and such a feature suggests that the conscious states of linguistic creatures are such that their subjects are aware of being in them. But it does not follow that *only* the conscious states of linguistic creatures are such their subjects are aware of being in them. Here I would advert to the symmetry argument, which applies to both linguistic and non-linguistic creatures. It is reasonable to suppose that dogs, dolphins, gorillas, etc. have phenomenally conscious experiences, i.e., have experiences that are like something. But then just as a physical object or event cannot be like something to a dog (dolphin, gorilla, etc.) unless the dog is aware of that object or event, so a conscious experience cannot be like something to a dog unless the dog is, in some way, aware of that experience.

The Symmetry Argument and the Argument from Spontaneous Reportability are not meant to suggest that the way we are aware of our experiences is akin to the way we are aware of ordinary physical objects and events. That is, I do not mean to suggest that we should adopt an "inner sense" model of self-awareness (see ch. 5), according to which the awareness of our own mental states can be accounted for on

analogy with sense perception. My claim, for the moment, is only that if a subject is conscious of *x*, then the subject is aware of her consciousness of *x*. A conscious state, in other words, is always a state of which the subject is, in some appropriate way, aware. I undertake the task of articulating the most plausible model for this view of state consciousness in succeeding chapters.

CHAPTER 5

Higher-Orderism

5.1 Preamble

Our main result from the previous chapter is that all conscious states necessarily involve self-awareness, in the sense that if a subject is in a conscious state, then the subject is, in some appropriate way, aware of being in that state. My aim in this chapter is to begin elaborating on the "in some appropriate way" clause of this claim. Recent analyses of consciousness that take conscious states to necessarily involve self-awareness are roughly divided between higher-order (HO) or "two-state" theories, according to which a mental state is conscious in virtue of a distinct second-order state directed toward it, and intrinsic or "one-state" theories, according to which the self-awareness involved in conscious states is an intrinsic feature of those states. In this chapter I turn our spotlight on the HO theory of consciousness. It will emerge in what follows that the HO theory is not without its attractions. A number of conscious phenomena can be effectively accounted for within the theory, and the theory is flexible enough to safely absorb a number of *prima facie* difficulties. Despite these attractions, however, the HO theory is, in the end, unsustainable: it rests on at least one implausible assumption and has at least one grossly counterintuitive consequence.

My purpose, however, is not merely to cast aspersions on the HO theory. As anyone acquainted with the pertinent literature well knows, objections to the theory, of which there are many,[1] have nowise daunted committed proponents of the view.

---

1. For further critical commentary on the HO theory, see: Aquila 1990; Baker 2000, ch. 3; Block 1995; Byrne 1997, 2004; Chalmers 1996, pp. 230–1; Dretske 1993, 1995, 1997; Droege 2003, ch. 2; Francescotti 1995; Goldman 1993; Güzeldere 1995; Hellie 2007a; Hossack 2002, §II; Levine 2001, ch. 4; Lurz 2003, 2004; Manson 2002; Natsoulas 1993; Neander 1998; Papineau 2002, ch. 7; W. Robinson 2004; Rowlands 2001, chs. 4 and 5; Schröder 2001; Seager 1999, ch. 3, 2004; Siewert 1998, ch. 4, §4.4 and ch. 6; Stubenberg 1998, esp. chs. 4, 5, and 8; Thomasson 2000; Woodruff Smith 1986, 1989, 2004b; Zahavi 1999, 2004; and Zahavi and Parnas 1998. See Gennaro 2003 for a defence of the HO theory against Papineau's criticisms; and see Jacob 1996 for a defence against Dretske's criticisms. Also see Francescotti (1995, pp. 243–5), Byrne (1997, pp. 113–14), and Seager (1999, pp. 65–72), who, although critical of the HO theory, argue that Dretske's main objection is unsuccessful.

In addition, the objections I expound are by no means original (though I arrived at the phenomenological objection – see §5.2.2 below – independently of its other purveyors), and so there is reason to think that HO theorists, assuming that they are familiar with the objections, will find them wholly unconvincing. However, the theory is worth examining for two reasons, the most important of which is that it will allow me to mark off the theory from Brentano's one-state theory of consciousness. While a glimpse of the contour lines of the HO theory suggests that it is similar to the Brentanian model – in that, among other things, it construes conscious states in terms of self-awareness – a more painstaking examination of the HO theory in this chapter, and of the Brentanian model in the next, will make it evident that these surface similarities disguise a number of profound dissimilarities, dissimilarities that go well beyond the obvious difference between the theories, namely, that the former, unlike the latter, is a two-state theory. Second, since I think a roughly Brentanian model of consciousness – i.e., a Brentanian model augmented with certain Sartrean considerations – is superior to the HO theory, examining the HO theory will afford me the opportunity to explain why I believe it ought to be rejected in favour of a roughly Brentanian/Sartrean account.

The chapter is structured as follows. In section 5.2 I provide a rather tedious exposition of a version of the HO theory that ought, in my judgement, to be considered the exemplary source, providing some signposts along the way that indicate how it is similar to, and how it differs from, the Brentanian model. Then, in sub-sections 5.2.1 and 5.2.2, I spell out two grave objections to the theory. This leaves the field open for a more careful examination of the Brentanian model, a task I undertake in the next chapter.

5.2 The Higher-Order theory of consciousness

Allegiance to some version or other of the HO theory of consciousness is noticeably prevalent in current philosophical discussions of consciousness. Among the theory's more prominent recent exponents are Armstrong (1968, 1981), Lycan (1996, 2001, 2004), Gennaro (e.g., 1996, 2005, 2006), and especially Rosenthal (e.g., 1986, 1990, 2004a). In addition, Nelkin (esp. 1989a, 1989b, 1993) has promulgated a view that resembles the HO theory, and Dennett (1993, p. 928) has remarked that the theory is "very close kin" to his own view of consciousness. Even Dretske, a professed adversary of the theory, makes gestures toward favouring some version of it.[2]

2. According to Dretske's representational theory, representational states must be made available to higher-level cognitive processes in order to qualify as conscious experiences; they must be given "conceptual form" (Dretske 1995, p. 20). This is very congenial to HO theory.

Rosenthal and Gennaro are advocates of what is properly called the higher-order thought (HOT) theory of consciousness, which is to be distinguished from the higher-order experience or perception (HOP) theory of consciousness, for which Armstrong and Lycan are leading spokesmen. The idea behind the HOP theory is that a mental state is conscious only if it is the object of a higher-order perception or experience. Lycan characterizes the HOP theory as an "inner sense" conception of conscious awareness. Influenced by the Lockean doctrine that "consciousness is the perception of what passes in a Man's own mind" (Locke 1975, bk. II, ch. 1, §19),[3] Lycan hypothesizes that consciousness is a "perception-like second-order representing of our own psychological states and events" (1996, p. 13). The basic idea behind the HOT theory, on the other hand, is that in order for a mental state to be conscious the subject must have an appropriate higher-order belief, thought, or judgement *that* he or she is in that mental state. The principal difference, then, between HOT theories and HOP theories is that HOT theories maintain, while HOP theories do not, that the higher-order state must be intellectual or cognitive in nature, involving some kind of conceptual component.

As I suggested, however, a core feature of HO theories, whether HOP or HOT, is that no mental state is conscious solely in virtue of its intrinsic characteristics. Both HOP and HOT theorists posit two *distinct* "levels" or "tiers" of mental activity, one of which monitors the other – quasi-perceptually or experientially in the case of HOP theories and intellectually or cognitively in the case of HOT theories – rendering some of its contents conscious. The following is a somewhat loose formalization of this idea:

> *Core thesis of HO theory*: a mental state M of a subject S is conscious at *t* only if at *t* M is the object of an appropriate mental state M_1, and $M \neq M_1$.

The guiding idea is that a mental state is rendered conscious by a numerically distinct higher-order thought or experience. Consciousness, according to HO theory, is an extrinsic, relational property of certain mental states, not an intrinsic property.[4]

---

3. Lycan (1996, p. 13) also cites Kant as an intellectual forerunner of the HOP theory, but Kant is, I think, more accurately construed as a forerunner of the HOT theory, especially in light of his general argument for the claim that it is necessary to apply *concepts* to experience. See Gennaro (esp. 1996, but also 2005) for a discussion of the connections between the HOT theory and Kant.

4. Though at least one HOT theorist, Carruthers (1996, 2000), takes consciousness to be an intrinsic property of certain mental states. According to Carruthers' dispositionalist version of the HOT theory, it is one and the same mental state that possesses both first-order non-conceptual content and higher-order non-conceptual content. In addition, Gennaro's version of the HOT theory posits an intrinsic relation between higher-order thoughts and the mental states they are about.

The idea that consciousness is an extrinsic property would, of course, play out rather differently in HOP and HOT theories, but the idea is reasonably clear. In what follows I will restrict myself to a discussion of the HOT theory, though many, indeed most, of my comments will pertain *mutatis mutandis* to the HOP theory.

HOT theories come in different varieties. According to Rosenthal and Gennaro's *actualist* versions of the theory, a mental state is conscious in virtue of its co-occurrence with a suitable higher-order thought. More precisely, a mental state is conscious only if the subject of the mental state has a roughly concurrent higher-order indexical-assertoric thought to the effect that he is in that state. Thus, on the actualist version of the HOT theory, I am consciously seeing the grass at *t* only if at *t* I have a higher-order thought to the effect that I am seeing the grass.[5] Rosenthal states the view as follows:

> Conscious states are simply mental states we are conscious of being in. And, in general, our being conscious of something is just a matter of our having a thought of some sort about it. Accordingly, it is natural to identify a mental state's being conscious with one's having a roughly contemporaneous thought that one is in that mental state. (1986, p. 335)

According to a *dispositionalist* version of the theory, by contrast, a mental state is conscious in virtue of its *availability* to a higher-order thought. On this view, "the HOTs [higher-order thoughts] which render [mental states] conscious are not necessarily actual, but potential" (Carruthers 2000, p. 227). In the rest if this chapter I will focus on the Rosenthal's well-developed actualist version of the HOT theory. Rosenthal's version of the theory is partly responsible for the widespread enthusiasm for the HOT theory, so his treatment bears careful consideration.[6]

One of the principal reasons Rosenthal favours a two-state model of consciousness is because he thinks that consciousness can be informatively explained only on the assumption that it is an extrinsic, relational property of certain mental states.

---

5. A higher-order thought, according to Rosenthal, is "any episodic intentional state with an assertoric mental attitude" (1993c, p. 913n2). And elsewhere he calls a higher-order thought an "assertoric, occurrent propositional state" (1990, p. 742). Higher-order thoughts are sometimes described, as well, as "occurrent beliefs" (e.g., Byrne 1997, p. 106), so to have a higher-order thought about being in a particular mental state is to have an occurrent belief about being in that mental state. This usage is odd, however, for it suggests, implausibly, that beliefs are mental states.

6. This decision to focus on an *actualist* version of the theory is not arbitrary. I find persuasive the claim that "the consciousness of mental states is phenomenologically something occurrent. Since consciousness does not appear to be dispositional, it is ad hoc simply to posit a disposition that comes and goes as needed" (Rosenthal 1986, p. 338).

Rosenthal reasons as follows. If consciousness can be informatively explained, then it must have some sort of articulable structure. But if consciousness is an intrinsic property, then it does not have an articulable structure, i.e., it is "simple and unanalyzable" (1990, p. 736). Therefore, if consciousness is an intrinsic property, then it cannot be informatively explained.[7] In the next chapter, I confront this argument as a possible objection to the Brentanian model of consciousness, so, for present purposes, I propose to grant that it is sound.

The HOT theory may be viewed as an attempt to explain the intransitive consciousness of mental states in terms the transitive consciousness of creatures (see Rosenthal 1990, pp. 737–8). Rosenthal proposes, accordingly, that a mental state M of a subject S is intransitively conscious only if S is transitively conscious of being in M, i.e., only if S has a higher-order thought about being in M. Explaining intransitive consciousness in terms of transitive consciousness is not circular, Rosenthal argues, because transitive consciousness and intransitive consciousness are distinct states: "transitive consciousness can occur without intransitive state consciousness.[8] One is transitively conscious of something if one is in a mental state whose content pertains to that thing – a thought about the thing, or a sensation of it. That mental state need not be a conscious state" (1990, p. 737). To illustrate, Rosenthal invites us to consider the "cocktail-party effect," which "occurs when one screens out the sounds of conversation other than one's own. Still, if one's name is mentioned in a conversation one had screened out, one's attention immediately shifts to that conversation, showing that one must have had an auditory consciousness of what was being said" (1990, p. 743). This example indicates, according to Rosenthal, that a subject can be conscious of *x* without the state in virtue of which the subject is conscious of *x* itself being conscious.

That a subject can be conscious of *x* (say a target mental state) without the mental state (the HOT) in virtue of which the subject is conscious of *x* itself being conscious also blocks the threat of an infinite regress of higher-order states. If the target mental state is conscious only if it is accompanied by a second-order *conscious* thought, then the consciousness of the second-order thought would have to be explained in terms of a conscious third-order thought. But then the consciousness of the third-order thought would have to be explained in terms of a conscious fourth-order thought, and so on *ad infinitum*. Some regresses need not be vicious,

---

7. Rosenthal advances an argument along these lines in a number of places: e.g., 1986, pp. 344–5; 1990, p. 736; 1993b, p. 157.

8. Though the converse does not hold: intransitive state consciousness, according to Rosenthal, always implies transitive creature consciousness, i.e., always implies that the subject of the intransitive state is aware of being in it (see Rosenthal 1990, pp. 743–4; 1993a, p. 359).

but this one would be, since, under the HOT theory, it is a necessary condition for the existence of conscious mental states that the regress be *completed*. Thus, if the regress is not completed, conscious mental states would not exist. But, of course, conscious mental states do exist. As such, the HOT theory states that *every* conscious mental state has, "at its top" as Rosenthal puts it (1986, p. 337), an unconscious HOT. Consequently, the HOTs that confer consciousness upon our mental states are rarely themselves conscious.[9] One would have a conscious thought only via a third-order thought about the second-order thought, which means that the initial HOT in the hierarchy of HOTs is always unconscious unless it is accompanied by a HOT about it, as in cases of introspection.[10]

The wording here (and elsewhere in this chapter), though Rosenthalian, is potentially misleading. I do not mean to suggest that the HOT's "conferring" consciousness on the mental state it is about (or "making" the mental state conscious, or being a "source" of consciousness) is a matter of the HOT's causing some kind of *intrinsic* change to the mental state. As Rosenthal (e.g., 2000b, §III) has stressed, being aware of a mental state, under the HOT theory, results in no change to that state's intrinsic properties. The idea is not that the presence of an appropriate HOT

9. See Rosenthal 1986, p. 336; 1990, pp. 742–3; 1993c, p. 913; 1993d, p. 205; and 2002a, p. 658.

10. Rosenthal's treatment of introspection is somewhat nebulous. According to Rosenthal, when one introspects a mental state, one's higher-order thought about that mental state becomes conscious (1990, pp. 745–6). But, so stated, this seems false. If I think about one of my mental states, if, for example, I judge that I am in pain, my thought about my pain is not unlike a thought I might have about some extra-mental object; I might, for example, judge that it is raining. Thus, it does not seem as though introspection can be identified with conscious HOTs in the way Rosenthal suggests. Let me enlarge on this complaint. HOTs are indexical-assertoric mental attitudes, so if I non-introspectively judge that it is raining, the unconscious higher-order thought accompanying this judgement will have the structure: ⟨I herewith judge that it is raining⟩. And this HOT will not itself be conscious unless I have an unconscious HOT about it. Likewise, if I judge that I am in pain, the HOT accompanying this judgement will have the structure: ⟨I herewith judge that I am in pain⟩. Now, according to Rosenthal, this HOT *will* be conscious, since, by Rosenthal's lights, when one thinks about a mental state, one engages in introspection. But why suppose that the HOT accompanying my judgement that I am in pain is conscious, whereas the HOT accompanying my judgement that it is raining is unconscious? What, exactly, is the difference between judging that it is raining and judging that one is in pain? There does not seem to be any substantive difference in *kind* here: one can, it seems, judge that one is in pain without having a conscious thought to the effect that one is judging that one is in pain. The moral here is that Rosenthal's account of introspection requires modification. His blanket statement that HOTs are conscious in cases of introspection seems inadequate, at least if one wants to hold, as Rosenthal does, that relatively simple thoughts about one's mental states, like the one considered above, are introspective states.

causes, or produces, or gives rise to a conscious mental state M, but rather that the compresence of the HOT and M *constitutes* M's being conscious.

Another important feature of the HOT theory is that the HOTs that confer consciousness on our mental states are not based on conscious inference or conscious observation. Rosenthal elucidates this component of the HOT theory as follows:

> When a mental state is conscious, one's awareness of it is, intuitively, immediate in some way. So we can stipulate that the contemporaneous thought one has is not mediated by any inference or perceptual input. We are then in a position to advance a useful, informative explanation of what makes conscious states conscious. Since a mental state is conscious if it is accompanied by a suitable higher-order thought, we can explain a mental state's being conscious by hypothesizing that the mental state itself causes that higher-order thought to occur. (1986, pp. 335–6)

In this passage Rosenthal claims that mental states cause the HOTs of which our mental states are objects, but he subsequently rescinded the causal requirement between first-order mental states and HOTs, arguing that it "is very likely too strong" (1990, p. 744; also see Rosenthal 1993b, p. 159). According to Rosenthal, our mental states might play a part in causing HOTs, but they need not be the principal causal factor. "[M]ere accompaniment," he argues, "is all one needs to spell out the proposal to explicate intransitive in terms of transitive consciousness" (1990, p. 744; also see Gennaro 1996, pp. 73–5). As far as I can tell, nothing in the discussion immediately to follow hangs on the causal requirement. The important point is that the link between lower-order mental states and higher-order thoughts is suitably unmediated, i.e., based on neither conscious inference nor conscious observation.[11] As the passage quoted above indicates, Rosenthal makes this stipulation in order to guarantee the intuitive immediacy of our awareness of our own conscious states, i.e., in order to rule out cases in which we become conscious of a particular mental state even though we would not regard that mental state as conscious. For example, it is possible, according to Rosenthal, to be persuaded that one is angry through someone else's testimony. One realizes, in such cases, that one is angry in the absence of

---

11. Rosenthal allows for *unconscious* inference or observation to mediate between HOTs and the mental states they are about. "Neither inference nor observation," he says, "would impede our sense of immediacy if the inference or observation were not itself conscious. So we must rule out only HOTs based on inference or observation of which we are conscious" (2002b, p. 244n23). Ultimately, however, we need not trace the higher-order thoughts that render our mental states conscious to any specific aetiology: "the aetiology of the HOT does not matter, only the appearance of aetiology. A state is conscious only if we are conscious of it in a way that seems spontaneous and noninferential. As long as it seems that way, it does not matter how it is caused" (Rosenthal 2002d, p. 721).

any conscious feeling of anger.[12] Rosenthal argues that it is even possible to infer from one's own behaviour that one is angry while experiencing no conscious feeling of anger (1990, p. 737). Other cases suggest themselves. For example, S might infer that she is depressed on the basis of her recent behaviour, or on the basis of her psychotherapist's testimony. But even if S really is depressed about something, her newfound awareness of her depression would not make it conscious.

With the non-inferentiality requirement between HOTs and lower-order mental states in hand, we can formulate the HOT theory as follows:

> *HOT theory*: A mental state M of a subject S is intransitively conscious at *t* if and only if at *t* S has a non-dispositional, noninferential higher-order assertoric thought about being in M.

At this point it is worth noting a point of convergence between the HOT theory and the Brentanian model, a point of convergence beneath which lies an important point of divergence. As I explain in detail in the next chapter, the Brentanian model also tries to account for the intuitive immediacy of one's awareness of one's conscious states; however, the Brentanian model, unlike the HOT theory, has an intrinsicality component that ensures this immediacy. To be more precise, one's awareness of one's conscious states, on the Brentanian model, is an intrinsic feature of those states, which means that the question of whether there is a need to stipulate that one's awareness of one's conscious states is non-inferential is moot, since if one's awareness of one's conscious states is an intrinsic feature of those states, then, *ipso facto*, one's awareness of one's conscious states is non-inferential. Thus, the problem of immediacy does not arise for a theory of state consciousness that takes the awareness of one's conscious states to be intrinsic to those states.[13]

---

12. There is a possible problem here that I will mention but not pursue. Someone unfamiliar with the niceties of the HOT theory might object to the theory on the grounds that it seems possible to have a HOT about being in a particular mental state even though we would not regard that mental state as conscious. Rosenthal would reply, of course, that a HOT renders a mental state conscious only if it is non-inferential. But this response seems unsatisfactory; after all, one might reasonably ask *why* the HOT must be non-inferential, and it seems arbitrary and ad hoc to reply that it must non-inferential because it is possible to have a HOT about being in a mental state even though we would not regard that mental state as conscious. The difficulty, in essence, is that, envisaging a difficulty, Rosenthal simply stipulates that conscious-rendering HOTs are non-inferential. But such a stipulative manoeuvre seems to paper over, rather than remove, the envisaged difficulty. For further critical remarks on this aspect of the HOT theory, see Kobes (1995).

13. Rosenthal concedes this point: "One way to answer the question of immediacy is just to stipulate that one's being conscious of a conscious mental state is internal to that state itself. Immediacy is this guaranteed. Our being conscious of the state would be a part or aspect of the state itself, so nothing could mediate between the state and one's being conscious of it" (1993b, p. 157).

With an eye toward revealing some of the HOT theory's additional complexities, I should like to anticipate and forestall an apparent difficulty with the theory. For a mental state to be conscious, according to the HOT theory, the subject must have a suitably unmediated higher-order thought about being in that state, but, as stated above, this higher-order thought will not itself be conscious unless the subject has another higher-order thought about *it*. So a significant number of higher-order thoughts, indeed most, are unconscious. But how, it might be asked, can an unconscious mental state, a HOT, operate to confer consciousness on another unconscious mental state? If the HOTs required for consciousness can themselves be unconscious, then the HOT theory owes us an explanation of how these *unconscious* mental states can be the source of *conscious* mental states.

Van Gulick (2005, 2006), Goldman (1993, p. 366), Woodruff Smith (1986, p. 150; 1989, pp. 84-6), Stubenberg (1998, ch. 8), Rowlands (2001, p. 301ff.), and Thomasson (2000, p. 198) have all expressed doubts about whether this problem – what Van Gulick calls the "generality problem" – has an adequate solution. In addition, some have argued that the problem leaves the HOT theory vulnerable to zombie objections (e.g., Goldman 1993, p. 366; Dennett 1991, p. 311). If the mental states required for consciousness are themselves unconscious, then we can imagine a creature that harbours these mental states yet remains entirely unconscious. A zombie objection in this context (as in any other) strikes me as unpersuasive, however, for it seems to amount to nothing more than the plainly vacuous claim that the HOT theory might be false. But there *is* a difficulty, for it is puzzling how an unconscious HOT can confer consciousness on another unconscious mental state. In fact, the difficulty, as Dretske has pointed out (1995, pp. 100–1), pertains to *conscious* HOTs too, for it is equally puzzling how a conscious HOT can confer consciousness on an unconscious mental state.[14]

Rosenthal is well aware of this difficulty, of course, but I must confess that I do not understand his rebuttal. The objection, he avers,

---

14. There are additional puzzles; for instance, the HOT theory seems to lack an explanation of why only mental states are conscious. To use Goldman's classic example, "A rock does not become conscious when someone has a belief about it. Why should a first-order psychological state become conscious simply by having a belief about it?" (1993, p. 366). Dretske (1995, p. 100) has propounded a similar example, arguing that I can become conscious that I have cancer, i.e., my doctor can tell me, in which case I might have a higher-order thought to the effect that I have cancer. But my having such a thought does not render my cancer conscious. Byrne (1997, pp. 110–11) has offered a compelling response to this difficulty, however. According to him, it is analytic that only mental states can be conscious, in the same way it is analytic that only males can be brothers.

> disregards the distinction between transitive and intransitive consciousness. HOTs confer intransitive consciousness on the mental states they are about because it is in virtue of those thoughts that we are transitively conscious of those mental states.
>
> (1990, p. 743)

As far as I can tell, however, this is just a circumlocutory way of restating the basic HOT thesis. Rosenthal seems to be saying that we are conscious of our mental states because of the higher-order thoughts we have, and this is what makes our mental states conscious. But the question at issue, i.e., the question of how an unconscious mental state can confer consciousness on another unconscious mental state, has not been answered.

Shortly after the previous quotation, Rosenthal says:

> A state of consciousness can be a conscious state, that is, a state one is conscious of being in. Or it can be a mental state in virtue of which one is conscious of something. Adapting our terminology, we can call these intransitive and transitive states of consciousness, respectively. For a mental state to be conscious is for it to be a state one is conscious of being in. So a HOT can be a source of consciousness for the mental state it is about because the HOT is a transitive state of consciousness; it does not also need to be an intransitive state of consciousness.
>
> (1990, p. 743)

Once again, however, this is scarcely informative. Rosenthal here appeals to the distinction between transitive consciousness and intransitive consciousness, arguing that since a higher-order thought is a transitive state of consciousness, i.e., a mental state in virtue of which one is conscious of something, it can be a source of intransitive consciousness. But the crucial question of how a higher-order thought, itself unconscious, can be a source of consciousness remains unanswered.

I suspect that, in the final analysis, defenders of the HOT theory will have to remain tentative about how unconscious HOTs confer consciousness upon the mental states they are about. Indeed, it seems difficult to conceive of *a priori* evidence that a HOT theorist could recruit in support of a specific answer to the question posed. But then there is no reason to conclude from this that the HOT theory is inherently unstable. Given the natural limits of *a priori* theorizing, any further progress might well require wading into empirical waters. In addition, we do not want to ask too much of the HOT theory. The HOT theory, as Rosenthal is keen to stress, is not a "scientific explanation" of state consciousness (2002b, p. 235),[15] but

15. Though, importantly, it is still an *empirical* theory, in the sense that it is "in principle empirically falsifiable" and "sometimes supported by empirical and scientific evidence" (Gennaro 2005, p. 15n17). But then, as Gennaro (ibid.) has rightly pointed out, this empirical aspect of the theory is consistent with the assumption that the basic truth of the theory is known *a priori*.

rather an account of the structure of conscious states: "What we want is not the mechanisms or laws governing consciousness, but an answer to the preliminary question of what it is for a mental state to be conscious in the first place" (ibid.). So to expect an answer from the HOT theory to the question of how unconscious HOTs physically (or naturalistically) "yield" consciousness is to expect an answer to a question that the HOT theory does profess to have the resources to answer.[16]

The HOT theory, as we have seen, is more multifaceted than one might initially suppose. What is more, in addition to the *prima facie* difficulty rehearsed above, the theory seems to have the resources to neutralize another familiar objection: the objection from infant and animal consciousness. This objection, which has been advanced by, among others, Dretske (1995, pp. 110–12), Lurz (2003), and Seager (2004) is *prima facie* compelling and, in Dretske's view, "decisive" (1995, p. 110), but the HOT theory can, I believe, deflect it. According to the HOT theory, we should identify a mental state's being conscious with one's having a roughly contemporaneous thought to the effect that one is in that state, but it is unclear whether infants and animals possess sufficiently rich conceptual resources for such higher-order thoughts. A mouse, for example, is doubtless able to smell cheese, but it is extremely unlikely that the mouse is capable of having the thought *that* it is smelling cheese. Human infants, similarly, experience pain, but it is doubtful that they have the ability to conceptually represent themselves *as* being in pain. Or so the objection goes. HOT theorists typically respond to this challenge, plausibly enough, by arguing that the HOT theory imposes relatively modest demands on infants and animals, since the HOTs required for consciousness are less sophisticated than one might think (e.g., Rosenthal 1990, pp. 741–2; Gennaro 1996, pp. 78–9).[17]

It seems, then, that the HOT theory can answer both the generality objection and the objection from infant and animal consciousness. However, as I mentioned in the preamble, the theory ultimately fails: at least two overwhelming objections, one technical and the other phenomenological, vitiate it. I will state the technical

---

16. Though see Gennaro (2005) for an interesting discussion of how a HOT theorist might go about answering this question.

17. Interestingly, Carruthers has disputed this response to the difficulty, arguing that the HOTs required for consciousness are a good deal more sophisticated that his fellow HOT theorists suppose (2000, pp. 193–6). This result does not lead Carruthers to conclude that higher-orderism is problematic, however. On the contrary, Carruthers stands by his dispositionalist version of the HOT theory and concludes that infants and animals do not have phenomenally conscious experiences. This is a most unhappy and counterintuitive conclusion, one that, to my mind, immediately casts doubt on any theory committed to it. Can anyone reasonably deny that there is something it is like for a teething infant to experience pain? For a response to Carruthers, see Gennaro (2004b, §1.2).

objection first, which is essentially a version of the objection from "targetless" HOTs as developed by, most notably, Byrne (1997), Neander (1998), and Seager (1999, pp. 72–84).[18]

5.2.1 Confabulated conscious states

As we have seen, Rosenthal believes that a condition of adequacy for an explanation of the nature of consciousness is that it yields a plausible account of the nature of intransitive state consciousness, and he thinks the best way to satisfy this desideratum is in terms of transitive creature consciousness. Now explaining intransitive state consciousness in terms of transitive creature consciousness non-circularly involves stipulating that transitive consciousness and intransitive consciousness are *distinct* states. But the HOT theory goes awry in its reliance on the supposed division between these states. If transitive state *k* (the HOT) and intransitive state M are distinct states, and if *k's* mental target need not be implicated in causing *k*, then *k* can occur without M and vice versa. If M occurs without *k*, then M is simply an unconscious mental state, but if *k* occurs without M, then the subject of M has what Rosenthal calls a "confabulated" conscious state (Rosenthal 1990, p. 744; 2000a, p. 211). Confabulated conscious states are, essentially, targetless HOTs; they occur when the subject has a thought to the effect that she is in a particular mental state even though the mental state does not exist. Despite Rosenthal's attempts to minimize their counterintuitive implications and even embrace them as explanatorily indispensable (see Rosenthal 2000a),[19] confabulated conscious states impale the HOT theory on the horns of an insuperable dilemma. On the one hand, it cannot say that S's confabulated conscious state is a genuine conscious state, since the requisite lower-order mental state does not exist. On the other hand, it is committed to saying that it *seems* to S as though she is in a conscious state, since having a HOT about being in a particular mental state, even a non-existent one, suffices to make the subject believe she is in that mental state. Indeed, there is no difference, to the subject, between a veridical experience of seeing *x* and a confabulated experience of seeing *x*. The upshot is that S thinks she is in a conscious state even though she is not. I take this to provide a formidable *prima facie* case against the HOT theory.

Suppose I have a higher-order thought to the effect that I am in pain in the absence of the appropriate lower-order mental state. According to the HOT theory, I will not actually be in a conscious pain state because the lower-order state does

18. Caston (2002, pp. 780–1), Levine (2001, pp. 108–9), Kriegel (2003b, pp. 119–20), and Rowlands (2001, pp. 92–3) have also pressed this objection.

19. See also Armstrong (1984, pp. 135–7) and Lycan (1996, pp. 17–23).

not exist. All the same, since I have a higher-order thought to the effect that I am in pain, I will feel pain and believe that I am in pain. Moreover, I may take action to alleviate the pain. The point here is that my confabulatory pain state will be subjectively indistinguishable from its veridical counterpart, and so I will feel (and perhaps even act) as though I am in an authentic pain state. Now the difficulty facing the HOT theory, as I expect is already plain, is that it is committed to upholding the conclusion that I am not consciously experiencing pain even though I am feeling pain, and perhaps actively seeking pain relief, which is incoherent. As I pointed out in Chapter One, it is part of the very concept of pain (and of our sensation concepts generally) that if it seems to one as though one is in pain (which it surely does if one feels it), then one is.

Notice that, at the risk of divesting the HOT theory of all its explanatory potential, one cannot respond to this difficulty by saying that confabulated conscious states are genuine conscious states with their own phenomenality. If a subject does not require a lower-order mental state to have a full-fledged conscious experience, then the HOT theory does not provide the sort of "two-state" explanation of consciousness it purports to provide after all, since higher-order thoughts bear the entire explanatory burden. That is to say, there is no need to postulate the existence of lower-order mental states as distinct from higher-order thoughts, since higher-order thoughts suffice for consciousness. But if the higher-order thoughts required for consciousness can make a creature conscious in the absence of the appropriate lower-order state, then the HOT theory merely recapitulates the problem of what constitutes a conscious state at the level of higher-order thoughts, for surely some account has to be given of what the consciousness of these targetless higher-order thoughts consists in.

Gennaro (e.g., 1996, 2006) has elaborated a version of the HOT theory that may have the resources to surmount the foregoing difficulty. On Gennaro's version, consciousness is an intrinsic feature of certain mental states. All conscious mental states must be accompanied by a higher-order thought – meta-psychological thought (MET) in Gennaro's terminology – but the MET that renders a mental state conscious is a *proper part* of that mental state, where a proper part is not identical with the whole of which it is a part. So conscious rendering METs are directed at the states (or at parts of the states) of which they are part.[20] Gennaro calls this the "wide intrinsicality view" or WIV:

---

20. In his 1986 Rosenthal considers and then rejects the idea that we can make sense of the claim that our higher-order thoughts are part of the mental states they render conscious. According to Rosenthal, "we have no nonarbitrary way to tell when one mental state is part of another. Accordingly, there is no reason to uphold the idea that our awareness of conscious states is a part of those states other than a desire to sustain the Cartesian contention that all mental states are conscious states. Moreover, if conscious states have parts in this way, the question arises whether

> We can understand consciousness as involving the property of 'accompanied by a MET' in much the same way as Rosenthal. But we might individuate conscious states 'widely,' i.e., in a way that treats consciousness as an intrinsic property of those states. On this account, the MET is part of the conscious state. I will call it the 'wide intrinsicality view,' or WIV. Conscious states are complex, and consciousness is an intrinsic property of conscious states.
>
> (1996, p. 23)

The WIV might be able to sidestep the problem of confabulated conscious states because although it holds that mental states can be unconscious, i.e., that they can occur in the absence of any conscious rendering MET, it does not treat conscious rendering METs as distinct from conscious states (Gennaro 1996, pp. 73, 74). The idea, as it see it, is that we have a number of unconscious mental states, some of which have a MET as proper part and which become conscious when stimulated or activated by that MET. But since the MET is part of the mental state it renders conscious, it cannot occur in the absence of that state. So mental states can occur in the absence of METs, but METs cannot occur in the absence of mental states. Gennaro puts it this way:

> What makes a mental state conscious is the presence of a MET directed at that state, but a MET cannot bring about a conscious mental state if the state does not exist in the first place. . . . If a system *has* a mental state M, then a MET directed at *it* will render it conscious. But M must already be present.
>
> (1996, p. 85)

It appears, then, that Gennaro's version of the HOT theory is immune to the problem of confabulated conscious states, since if a MET cannot occur in the absence of the mental state it purports to be about, then the MET cannot cause a confabulated conscious state.

Appearances are deceiving. The problem of confabulated conscious states bedevils Gennaro's theory at the level of *introspective* METs. Introspection, on Gennaro's theory, is a complex form of self-awareness that involves having *conscious* METs. Now a MET is conscious when it is accompanied by a MET about *it*, which means that introspection involves having three mental states: (1) a first-order mental state, (2) a MET directed at the first-order mental state, and (3) a MET, or MET', directed at the MET that is directed at the first-order mental state (see Gennaro 1996, pp. 24–5). For simplicity, Gennaro refers to the combination of a mental state and a MET as a *complex* state, so he refers to the introspective combination of a MET

---

all the parts of such states must be conscious, or only some. If all, then the awareness of the mental state will have to be conscious. A regress would thus arise that is exactly parallel to that which arose when we construed the awareness of conscious mental states as due to a distinct higher-order thought" (1986, p. 345). See Gennaro (1996, p. 26) for a response to this argument.

and a MET' as a complex MET. Thus, introspection involves having a first-order mental state accompanied by a complex MET. Now since all conscious states, on Gennaro's view, are mental states that are accompanied by a MET, Gennaro must hold that all complex states are, in and of themselves, conscious states, since complex states *are* mental states that are accompanied by a MET. And since a complex MET is a complex state, a complex MET is, in and of itself, a conscious state. But since a complex MET is a conscious state, it cannot be part of the first-order state at which it is directed, since a conscious state cannot be part of a first-order state (nor can a first-order state that is not rendered conscious be part of a conscious state). METs are part of the first-order states at which they are directed, but once a MET and a first-order state suitably co-occur, they form a mental structure that constitutes a conscious state. Thus, conscious states are complex states that contain parts, but they are never part of other mental states, either first-order or second-order. This means, however, that when one engages in introspection, one's original MET detaches itself, so to speak, from the first-order state at which it is directed and combines with another MET to form a complex MET. But once this uncoupling occurs, the complex MET becomes distinct from the first-order state it purports to be about, which means it can occur, or at least persist, in the absence of that (or any other) first-order state. It follows that the subject of the complex MET could be in a confabulated introspective state, i.e., be reflecting on or thinking about an occurrent mental state (e.g., a pain) that he is not actually in.

But this poses a problem for Gennaro's theory. The problem is not that the subject could be reflecting on an occurrent mental state, say a pain, that he cannot recognize as a pain. There is nothing controversial about that: surely it is possible for a subject to experience pain, reflect on it, and yet fail to know what pain is, just as it is possible for a subject to experience an unfamiliar emotion, reflect on it, and yet fail to know what the emotion is. The problem, rather, is that the subject could believe that he is experiencing *some* sensation – perhaps the one we call "pain" – reflect on it, but nevertheless fail to be experiencing *any* sensation. But this, again, is incoherent, for it entails, absurdly, that it makes sense that a subject could sincerely withhold assent from the claim that he is experiencing some sensation at time *t* even though he sincerely avows experiencing that sensation at *t* (on the grounds that it seems to him as though he is) and is attending to it at *t*.

5.2.2 A phenomenological objection

Even if (some version of) the HOT theory can somehow deflect the objection from confabulated conscious states, phenomenology reveals a much less elaborate, yet much more potent objection to the theory, an objection to which all versions of the theory, including Carruthers and Gennaro's intrinsic versions, are equally

vulnerable. The difficulty stems from the fact that on all versions of the HOT theory one's higher-order thoughts are, more often than not, unconscious. This is because only in relatively rare instances, i.e., in cases of introspection, does one have conscious thoughts about one's own conscious states. By and large, when one is conscious of *x*, one does not have a simultaneous conscious thought to the effect that one's conscious experience of *x* is occurring. For example, when one sees *x*, one does not consciously think: ⟨I am now seeing *x*⟩ (or ⟨I now see *x*⟩);[21] when one remembers *x*, one does not consciously think: ⟨I am now remembering *x*⟩; and so on. But insofar as we seldom have such thoughts consciously, the HOT theory must affirm that the higher-order states in virtue of which we are conscious of our lower-order states are, in general, unconscious.

A few moments' reflection, however, reveals that this picture of the structure of conscious states is unacceptable on phenomenological grounds. It is accurate to affirm, as the HOT theory does, that normally we do not experience *higher-order* mental states of thinking about or perceiving our conscious states; but, to state a point shortly to be elucidated (ch. 6), we are nevertheless *experientially* aware of our conscious states.[22] My visual consciousness of the mountainous expanse in the distance, for example – even if I were not (as I am now) thinking about it – is accompanied by an experiential awareness of my visual consciousness of the mountainous expanse. Woodruff Smith describes this feature of consciousness in terms of the way subjects "*experience* consciousness" (1988, p. 30; 2004b, pp. 82, 94–5); and Zahavi describes it, along similar lines, in terms of one's having "experiential access to one's own consciousness" (2004, p. 78; 1999, pp. 21–4).[23] However, insofar as all HOT theories, including Carruthers and Gennaro's intrinsic versions, construe the awareness of our conscious states in terms of our having unconscious higher-order thoughts, they cannot capture this phenomenological fact.

There are three possible rejoinders to this objection, none of which succeeds in salvaging the HOT theory. First, one might simply deny, as I have simply affirmed, that one's awareness of one's mental states is conscious. I have made a phenomenological observation in maintaining that such self-awareness is an experiential phenomenon, but an objector might claim that she cannot find such experiential

---

21. Where the text between "<" and ">" represents the indexical-assertoric content of one's higher-order thought.

22. It is no surprise, therefore, that the view that all consciousness involves self-awareness, where this latter is *something we experience*, is almost universally endorsed, in one version or another, within the phenomenological tradition.

23. Also see Kriegel (2003b, pp. 127–8), who argues that, just as we dance dances and smile smiles, we experience our experiences.

awareness in her phenomenology. Does the objection thereby boil down to a phenomenological observation that a HOT theorist can reasonably reject merely by dint of a phenomenological counter-observation? Not at all. The chief motivation for theories of consciousness that take consciousness to lie in the awareness of one's own mental states is that there is an *intuitive* connection between consciousness and self-awareness. Rosenthal is explicit on this: "A strong intuitive connection obtains between being in a conscious state and being conscious of oneself" (1990, p. 744).[24] But whence this intuition? It is doubtful that it is a product of a theoretical investigation into the nature of consciousness, since those who take consciousness to lie in the awareness of one's own mental states have a *pre-theoretic* intuition that consciousness involves such self-awareness; they *come to the table* with it, so to speak. They do not explore the explanatory merits of different theories of consciousness and, as a result, *arrive* at the intuition. Quite the reverse: they have the intuition and subsequently formulate theories of consciousness that are based on it. It is therefore more reasonable to assume that the intuition is a product of *experience*, i.e., that we intuit experientially the self-awareness involved in conscious states. If this is correct, then the HOT theory simply disregards, by its own logic, the very phenomenon that motivated its construction in the first place.

Now Gennaro (2006, §3.4) has responded to the phenomenological objection to the HOT theory, and he is surely right to insist, contra Kriegel, that it is implausible that the *only* reason to posit such self-awareness is on first-personal, experiential grounds. Kriegel has claimed that "if we are inclined to admit such awareness *at all*, it is on first-personal, experiential grounds" (2003b, p. 121, my emphasis), which is false: there are also non-experiential, purely theoretical reasons for doing do. But this in no way detracts from the claim that such self-awareness is an experiential phenomenon, for even though there may be non-experiential reasons for positing it, the *initial* ground is always that it is phenomenologically revealed.

A second possible rejoinder to the phenomenological objection is to argue that the initial HOT in the hierarchy of HOTs is always conscious, so that one's lower-order mental state is rendered conscious by a conscious HOT. There are two obvious difficulties with this proposal. First, it does nothing to render the HOT theory any less phenomenologically idle. As I mentioned, when one is conscious of *x*, one rarely has a simultaneous conscious thought to the effect that one's conscious experience of *x* is occurring. Second, it results in the same vicious regress described earlier. If one's lower-order mental state is conscious if and only if it is accompanied by a *conscious* HOT then the consciousness of the second-order

24. Also see Rosenthal (1986, pp. 343, 345; 2002a, p. 234; 2004a, p. 17) and Kriegel (2004, p. 198); and cf. Lycan (2001).

thought would have to be explained in terms of a conscious third-order thought. But then the consciousness of the third-order thought would have to be explained in terms of a conscious fourth-order thought, and so on *ad infinitum*.

Notice that this regress cannot be parried by stipulating that the second HOT in the hierarchy of HOTs is unconscious. According to this suggestion, when one is conscious of *x*, the lower-order mental state, M, in virtue of which one is conscious of *x* is rendered conscious by a conscious HOT, H_1, and H_1 is rendered conscious by an unconscious HOT, H_2. The emerging position is this:

> *HOT theory (revised)*: A mental state M of a subject S is conscious at *t* if and only if at *t* S has a *conscious* non-inferential higher-order thought H_1 about being in M, such that H_1 is rendered conscious at *t* by an *unconscious* non-inferential higher-order thought H_2 about being in (having) H_1.

But surely this is ad hoc: one cannot simply legislate, without further ado, that a lower-order state is rendered conscious by a conscious HOT, whereas the conscious HOT is rendered conscious by an unconscious HOT.

The third response attempts to capitalize on the claim that an implicit self-awareness accompanies every conscious state. It might be suggested that the HOT theory can be revised to incorporate the insight. It is true that we rarely have conscious HOTs *explicitly*, but perhaps we have them *implicitly* (where such implicit awareness is not to be construed, as it often is in psychology and cognitive science, as an *unconscious* awareness). As considerations from the next chapter will make clear, however, adopting this strategy would militate sharply against the general dictates of the HOT theory, requiring such dramatic revisions to the theory that the resulting view would bear little, if any, resemblance to higher-orderism. Such a strategy would, for example, involve abandoning one of the HOT theory's fundamental tenets, namely, that lower-order mental states and higher-order thoughts are numerically distinct states.

To summarize, then, even though the HOT theory has the antecedent advantage of exploiting the assumption that consciousness is explicable in terms of self-awareness, the objections just reviewed are fatal to the theory, in my judgment. In the next chapter I expound the Brentanian model of consciousness, which shows promise of being able to avoid these difficulties while retaining the assumption that conscious states necessarily involve self-awareness.

CHAPTER 6

A "one-state" alternative

6.1 Preamble

It was agued in Chapter Five that, though it grasps some measure of the truth, the HO theory is ultimately unsuccessful. So if we want to preserve, as I think we must, the idea that conscious states necessarily involve self-awareness, then we will have to look to a different model. In this chapter I assess the viability of a "one-state" model of state consciousness of the kind Brentano advocated. There is reason to suppose that the Brentanian model is well-motivated. In addition to possessing both logical *and* phenomenological plausibility, it secures the explanatory advantages that HO theorists attribute to the HO theory – e.g., it provides an informative explanation of the structure of conscious states – while avoiding the HO theory's drawbacks.

Despite its strengths, however, the Brentanian model requires revision, since it implausibly construes the self-awareness involved in conscious states along "subject-object" lines, i.e., in terms of conscious states being objects to which subjects stand in a certain kind of relation. Later in this chapter, therefore, I augment Brentano's account with considerations from Sartre, who provided a non-objectualist account of the structure of state consciousness.

Let me state at the outset that the ensuing commentary is not intended to be an exercise in Brentanian (or Sartrean) exegesis. While I think my exposition of Brentano's theory is largely accurate, I occasionally take liberties with Brentano's arguments and recast them in a way in which Brentano scholars might not approve. In addition, I say nothing about Brentano's well-known opposition to materialism.[1] Nonetheless, even though the model of state consciousness expounded and defended in what follows is a reconstruction of Brentano's theory as I see it, it is still a reconstruction of Brentano's theory, and so I will call it "the Brentanian model."

---

1. Brentano held, famously (notoriously?), that intentionality cannot be explained in materialist terms. See his 1874, bk. II, ch. I. As far as I can tell, however, the Brentanian model is not irretrievably wedded to immaterialism, so one can, if one prefers, safely decouple the Brentanian model (at least as I expound it below) from Brentano's anti-materialist orientation.

6.2 The Brentanian model

Interestingly, Brentano is sometimes cited as a proto-HO theorist of sorts,[2] but Brentano's theory should actually be viewed as anathema to higher-orderism, at least if higher-orderism is understood as a two-state model of state consciousness of the kind explicated in the preceding chapter.[3] After all, Brentano wholly rejected the existence of *distinct* levels of mental activity, developing, instead, a model of state consciousness according to which consciousness is an intrinsic and immediate feature of certain mental acts,[4] not an extrinsic property conferred on a mental act from without. Furthermore, Brentano rejected the existence of unconscious mental activity, which most, perhaps all, higher-order theorists accept unequivocally. Brentano recognized that it is tempting to posit the existence of unconscious mental activity to block the regress that threatens accounts of consciousness that take consciousness to lie in the awareness of one's own mental states, but, as we shall see, Brentano devised a way to avoid the regress problem without making recourse to unconscious mental states.[5]

Citing Aristotle as his inspiration, Brentano argued that all conscious states are reflexive.[6] More fully, he argued that each mental phenomenon or mental act

---

2. See Gennaro 1996, pp. 27–9; Güzeldere 1995, p. 800n1; Rosenthal 1990, pp. 746, 752n60; 1993d, p. 211; 2002, p. 250; Siewert 1998, p. 357n1; and Zahavi 2004, p. 75.

3. Sartre has been portrayed as a proto-HOT theorist as well (see Gennaro 2002, and Siewert 1998, p. 357n1), but, for reasons I expound below, it is a mistake to suppose that Sartre espoused any kind of higher-order view.

4. Giving a satisfactory definition of intrinsicality is, it turns out, extraordinarily difficult, so difficult, in fact, that some (e.g., Dennett 1988, p. 635) have suggested that we abandon the notion of an intrinsic property altogether. But this seems rash. A central insight of the later Wittgenstein is that many of our concepts are imprecise, even hopelessly so, but this no more detracts from their usefulness in the commerce of everyday life than the having of three legs instead of four and wobbling a bit, detracts from the usefulness of a table. In what follows, I will rely upon a highly generic characterization, standard in the literature, according to which an intrinsic property is a non-relational property: *k* is an intrinsic property of *F* if *F's* having *k* does not rely on *F* bearing some relation to some other entity. (A useful discussion of the various subtleties of the intrinsic/extrinsic distinction can be found in Humberstone 1996).

5. Though, as I argue later in this chapter, one can accept Brentano's account of state consciousness without rejecting unconscious mental states.

6. See Caston 2002 for an excellent, probing commentary on Aristotle's reflexive approach to consciousness. Aristotle's main argument for his reflexive approach, i.e., the regress argument, can be found in the opening paragraph of his *De Anima* III.2 (425b11–17). I discuss the regress argument below.

(or psychical act)[7] is directed toward something, its *primary* object (or content[8]), but is at the same time directed toward itself, its *secondary* object.[9] "Every consciousness," he says, "upon whatever object it is primarily directed, is concomitantly directed upon itself" (1982, p. 25).[10] If I see a blue sky, the primary object of my act of seeing is the blue sky. But, in being so conscious of the blue sky, I am also aware of my act of seeing the blue sky. Thus, my act of seeing is a reflexive act or a form of self-consciousness. Brentano explains it so:

> Every mental act. . .includes within it a consciousness of itself. Therefore, every mental act, no matter how simple, has a double object, a primary and a secondary object. The simplest act, for example, the act of hearing, has as its primary object the sound, and for its secondary object, itself, the mental phenomenon in which the sound is heard.
>
> (1874, pp. 153–4)

Now, in most cases, when we see something (hear something, etc.) we are not introspectively aware, either casually or attentively, of our act of seeing. When I see something I am typically in the throes, so to speak, of seeing; I am visually conscious of whatever I see. I can, if I wish, reflect on or introspect my act of seeing, but in the usual course of things such spontaneous introspection is rare. But, according to Brentano, unintrospected mental acts no less involve self-awareness. In being concomitantly aware of my act of seeing, I do not reflect on my act of

---

7. Some have expressed misgivings about whether we can make sense of Brentano's use of the phrase "mental act" (e.g., Rosenthal 1990, p. 746). But the phrase need not cause concern, for Brentano understands it to mean the same thing as "mental state" (see, e.g., 1874, pp. 12, 155). In my exposition of the Brentanian model, then, I will use the expressions interchangeably. However, even if Brentano did intend "mental act" to refer to a mental performance of some kind, i.e., something we do mentally, his using the phrases "mental act" and "mental state" synonymously would not be inappropriate. After all, if one is concentrating (an act), then one is in a state of concentration; if one is listening to the sound of the ocean (an act), then one is in an auditory state; if one is imagining what it is like to be a bat (an act), then one is in a state of imagination; and so on.

8. Brentano used the term "content" to mean the same thing "object" (see, e.g., 1874, pp. 88, 100, 123, 138). Traditionally, the content of an experience cannot be identified with the object, but as far as I can see, nothing in the discussion to follow will be affected if I use the terms interchangeably.

9. It should be underscored that Brentano used the term "consciousness" as synonymous with "mental phenomenon" or "mental act" (see Brentano 1874, p. 102). So in my exposition of the Brentanian model, any occurrence of "mental act" (I will not use "mental phenomenon") can be replaced with "conscious mental act," "conscious mental state," or, simply, "conscious state."

10. The reference here is to a series of lectures given by Brentano between 1887 and 1891, published under the title *Descriptive Psychology*.

seeing, but I am nevertheless *implicitly* or *secondarily* aware of my act of seeing. When I see an object, there is always an implicit awareness of my act of seeing that accompanies my act of seeing. My act of seeing has, "on the side" as Brentano puts it, "itself as object" (1982, p. 26).[11]

It is important to distinguish very carefully between introspection and the sort of implicit self-awareness characteristic of the Brentanian model, since commentators sometimes conflate the two.[12] Acts of introspection are deliberate, moderately effortful, and generally unprotracted, whereas implicit self-awareness is a constant, effortless, and non-deliberate element of all conscious states.[13] One cannot voluntarily interrupt or halt one's implicit awareness of one's act of seeing, hearing, smelling.[14] Further, and relatedly, implicit self-awareness is a necessary condition for conscious states, whereas introspection is not. One can (and usually does) have unintrospected conscious states, but there is no state consciousness whatever without implicit self-awareness. In consciously seeing *x* (hearing *x*, etc.) a subject is *necessarily* implicitly aware of seeing *x*.[15]

---

11. To state a point I enlarge on in a subsequent section (§6.2.3), conscious states, on the Brentanian model, are considerably more complex than this characterization suggests, i.e., they possess what might be called *de se* content, or content that is directed upon *oneself*. Since nothing in this section hangs on characterizing conscious states in terms of their *de se* content, however, I will refrain, for the sake of brevity, from characterizing them in those terms.

12. See, for example, Searle (1992, pp. 141–3) and S. White (1987). Searle takes a mere three pages to dispense with the view that conscious states necessarily involve self-awareness, but his argument erroneously conflates implicit self-awareness and introspection. And White attempts to draw a fundamental distinction between conscious states and self-awareness, but his argument incorrectly assumes that being self-aware depends on the ability to engage in first person self-reference.

13. Also see Gurwitsch (1950, p. 6) and Woodruff Smith (1989, ch. 2, §3.2).

14. As Gurwitsch says: "At every moment, we have the possibility of reflecting upon and apprehending the act which we experience. But whether or not we realize [it] this possibility depends entirely upon our free choice, whereas we cannot help being aware of the act which we actually experience" (1950, p. 6).

15. Brentano does not argue, at least not directly, for the Transitivity Principle, i.e., for the principle that a mental state is conscious only if the subject is aware of it; rather, he simply assumes, by the definition of "conscious state," that this is the case. (His animadversions on unconscious mentality are revealing: he argues against the existence of unconscious mental states, but an unconscious mental state, for Brentano, is just a mental state of which the subject is not aware.) He then proceeds to offer a model of state consciousness predicated on the principle that he thinks honours both our intuitions and the merits of the relevant arguments. I invite those who are wary of the principle, however, to revisit the arguments propounded in its favour in Chapter Four (§4.4). See also below (§6.2.1).

It bears underscoring, however, that the Brentanian model, at least on my construal of it, does not *rule out* introspection, at least if introspection is understood, unobjectionably, as the act of thinking, casually or attentively, about one's own conscious experiences.[16] It is a relatively mundane fact that occasionally we do take up an explicit position vis-à-vis our conscious states.[17] I might, for example, become somatically conscious of the painful sensation in my foot, and subsequently reflect on the waxing and waning of its intensity; or I might shift my attention from the object of my act of looking to my act of looking itself, perhaps wondering why my vision is blurry. In such cases, a psychological state is the primary object of an act of reflection or thinking. But, in being so conscious of one's psychological state, one is implicitly aware of one's consciousness of that psychological state; one is implicitly aware of one's act of reflection. Brentano puts it this way:

> We can observe the sounds we hear, but we cannot observe our hearing of the sounds, for the hearing itself is only apprehended concomitantly in the hearing of sounds. On the other hand, when we recall a previous act of hearing, we turn toward it as a primary object, and thus we sometimes turn toward it as observers. In this case, our act of remembering is the mental phenomenon which can be apprehended only secondarily. The same holds true with regard to the perception of all other mental phenomena.
>
> (1874, pp. 128–9)

This consideration explains why one's concomitant apprehension of one's conscious states is such an elusive component of those states. Implicit self-awareness is a necessary condition for conscious states, which means that one's act of reflection (act of remembering, etc.) is not conscious unless it is accompanied by an implicit self-awareness. But then it follows that if one tries to reflect on one's implicit awareness of a particular conscious state, one's implicit awareness "shifts" to one's act of reflection. The implication: one cannot attend to one's implicit awareness of one's *occurrent* experience of *x*.[18]

The terminology being used here might occasion some perplexity. Brentano sometimes uses "secondarily" or, following Aristotle, "incidentally" to refer to

---

16. Though Brentano, justifiably I think, rejects introspection where introspection is understood as inner *observation* (1874, pp. 29–30). Brentano does not think it is possible for one to actually observe a mental act in the same way one can observe an extra-mental object.

17. In fact, as Merleau-Ponty has mentioned (1962, pp. 41, 47), to understand our consciousness itself it is *necessary* to "step away" from it and make it into an object of introspective consciousness.

18. Though this does not imply that one cannot attend to one's implicit awareness of one's *non-occurrent* experience of *x*. For more on this point, see Chapter Seven (§7.2.1).

one's concomitant consciousness of one's conscious states,[19] but, for purely stylistic reasons, I prefer "implicitly," so that is the term I will most often use. These terms are meant to signify that a conscious awareness of seeing, hearing, thinking, etc. is contained in the nature of any conscious act of seeing, hearing, thinking, etc., but is not directly expressed. When I see *x* (hear *x*, etc.), I am not reflectively or explicitly aware of seeing *x* – that is, I do not attend to or focus on seeing *x* – but I am nevertheless at least non-reflectively or implicitly aware of seeing *x*. As I argued in Chapter Four (§4.4.2), our ability to report on and describe our conscious experiences, even unintrospected ones, is evidence for this claim. I could, if prompted, spontaneously report that I was looking at *x*, but then this seems possible only if I was, at least implicitly, aware of myself looking at *x*.

Much of the relevant literature in psychology and cognitive science, and some in philosophy (e.g., Gennaro 2002), takes the term "implicit" and its cognate clan to be synonymous with "unconscious," so the possibility of confusion looms heavily here. Let me stress that the claim is not that the mental state in virtue of which one is aware of one's conscious state is an unconscious state. On the contrary, the mental state in virtue of which one is aware of one's state of seeing is itself a conscious state, albeit one that the subject has, as Brentano says, "on the side."

Now, as I previously indicated, perhaps the most crucial feature of the Brentanian model, at least for our purposes, is that one's act of hearing (act of seeing, etc.) and one's concomitant awareness of one's act of hearing are *non-contingently* or *essentially* related, in the sense that one's concomitant awareness of one's act of hearing is an intrinsic component of one's act of hearing. Put otherwise, one's act of hearing and one's concomitant awareness of one's act of hearing form a single mental act. To quote Brentano:

> The presentation[20] of the sound and the presentation of the presentation of the sound form a single mental phenomenon. . . . In the same mental phenomenon

---

19. Brentano writes: "The act of hearing appears to be directed toward sound in the most proper sense of the term, and because of this it seems to apprehend itself incidentally and as something additional" (1874, p. 128). And in the appendix to *Psychology from an Empirical Standpoint*, he writes: "When we said that reference to something as object is that which is most characteristic of mental activity, this should not be interpreted as though 'mental activity' and 'reference to something as object' mean exactly the same thing. Just the opposite is already clearly apparent from what we have said about every mental activity relating to itself as object, not, however, primarily, but secondarily or, as Aristotle, by whom the fact had already been noticed, puts it, 'incidentally' ('nebenbei')" (1874, pp. 275–6).

20. In Brentano's terminology, as may be recalled, to have a presentation is merely to be aware of an object and not make a judgement, take a stance, or otherwise have an attitude toward that object. See Brentano 1874, pp. 78–9.

> in which the sound is present to our minds we simultaneously apprehend the mental phenomenon itself.
>
> (1874, p. 127)

Later Brentano says:

> The consciousness of the primary object and the consciousness of the secondary object are not each a distinct phenomenon but two aspects of one and the same unitary phenomenon.[21]
>
> (1874, p. 155)

On the Brentanian model, then, consciousness is not an extrinsic property conferred on a mental act from without (e.g., by a numerically distinct, unconscious higher-order thought or perception). Rather, the features that make a mental act conscious are located within the act itself. Each mental act contains a consciousness of the primary object and a concomitant awareness of the secondary object, the act itself.

Appropriating some terminology from Caston (2002) and Natsoulas (2001) might be useful in shedding some light on this view. On Caston's interpretation of the Aristotle-Brentano line of argument, when one looks at a blue sky, there is a single *token* state that instantiates two *types* of activity, the activity of seeing a blue sky and the activity of being implicitly aware of seeing a blue sky.[22] But these activities do not, as it were, occur alongside each other; that is, they are not, in Natsoulas's words,

> so much constituents of a state of consciousness as they are features of it. When a state of consciousness occurs that belongs to a certain stream, the only mental act that can rightly be said to occur in that stream at the time is only that state of consciousness. This mental act is, for example, a visual perceptual awareness and also an awareness of itself as such.
>
> (2001, p. 228)

Put another way, a separate mental occurrence is unnecessary for there to be first-hand implicit awareness of one's own conscious state. One's conscious state and one's implicit awareness of one's being in that state – these two types of activity – are features of a single mental act.

The foregoing preliminary sketch of the Brentanian model reveals three points that are of cardinal importance: (1) a mental state M_1 of a subject S is conscious at *t* only if at *t* S has a mental state M_2, such that S is aware of M_1 in virtue of being in M_2; (2) the mental state M_2 in virtue of which S is aware of M_1 is conscious; and

---

21. I.e., one and the same unitary act, as the translators of Brentano 1874 (p. 155n2) point out.

22. There is nothing amiss about a token state or event instantiating two types. A red circle, to use one of Caston's examples, "is both coloured and shaped, a single token falling under two types" (2002, p. 794n96).

(3) M_1 and M_2 are not distinct, in the sense that it is impossible to have one without the other. These latter two points cut to the heart of the division between the Brentanian model and contemporary HO theories of consciousness. According to HO theories, M_1 and M_2 are distinct, logically independent states, which means that it is possible for one to occur in the absence of the other. On the Brentanian model, by contrast, M_1 and M_2 are not distinct, which means that it is impossible for one to occur in the absence of the other. In addition, according to the HOT theory, the higher-order states in virtue of which a subject is aware of her lower-order states are typically unconscious, whereas, on the Brentanian model, the "higher-order" states are always (at least implicitly) conscious.

6.2.1 The regress argument

Brentano's main argument for his model of consciousness, i.e., the regress argument, is derived from Aristotle, who, in his *De Anima* III.2, writes:

> Since we perceive that we see and hear, it is necessarily either by means of the seeing that one perceives that one sees or by another [perception]. But the same [perception] will be both of the seeing and of the colour that underlies it, with the result that either two [perceptions] will be of the same thing, or it [sc. the perception] will be of itself. Further, if the perception of seeing is a different [perception], either this will proceed to infinity or some [perception] will be of itself; so that we ought to posit this in the first instance.
>
> (425b12–17, as quoted in Caston 2002, pp. 769, 773[23])

Brentano employs this line of reasoning in the second book of his classic 1874 work, *Psychology from an Empirical Standpoint* (esp. pp. 121–30), frequently expressing his debt to Aristotle. The argument, put simply, claims that, since every perception is itself perceived, it is perceived either by itself or by another perception. But it cannot be perceived by another perception, since that would lead to an infinite regress of perceptions. Thus, it is perceived by itself. The regress, as Brentano points out (1874, p. 122), would be particularly vicious, since, in addition to the infinite iteration of perceptions, each perception would contain, within itself, an infinite complication of perceptions. This is because if one's visual perception of a blue sky – call it P_1 – were the object of a distinct perception, P_2, the blue sky would

23. Modern translators, Caston notes, typically render this passage in terms of *capacities* rather than *activities*. The Greek word that Aristotle uses, αιςθηςιο, which is usually interpreted as "sense," is ambiguous between: "a. the capacity. . . of perception, i.e., the sense [and] b. the activity. . . of this capacity, that is, the perception or perceiving" (Caston, 2002, p. 762). But Brentano, rightly according to Caston, renders the passage in terms of activities. In what follows I keep to the activity reading.

be presented twice, first as an object of P_1, and second as an object of P_2, whereas P_1 would be presented once (as an object of P_2). But since P_2 is the object of a distinct perception, P_3, the blue sky would be presented thrice, first as an object of P_1, second as an object of P_2, and third as an object of P_3, whereas P_1 would be presented twice (as an object of P_2 and P_3). And so on and so forth. Brentano, like Aristotle, focuses on perceptual experience, but his argument is meant to extend to other conscious states. The basic idea is that in order to avoid an infinite regress of conscious states without postulating unconscious states, one must assume that conscious states are reflexive.

As it stands, however, Brentano's argument is problematic, since it relies on an assumption that is now widely rejected, namely, that consciousness and mentality are coextensive. But the falsity of the assumption that all mental states are conscious implies the falsity of the claim that there is a threat of an infinite regress, since if it is not the case that all mental states are conscious, then one can block the would-be regress by holding that, at some (not-very-high) level, there is an unconscious mental state in virtue of which a subject is aware of her mental states.

In fact, however, as Thomasson has pointed out (2000, pp. 197–8),[24] there is no good reason why the existence of unconscious mental states should pose a problem for Brentano's thesis, even if it poses a problem for Brentano's specific argument for the thesis. Interestingly, Brentano does not rule out *a priori* the possibility of unconscious mental states. He says: "A presentation of the sound without a presentation of the act of hearing would not be inconceivable, at least *a priori*" (1874, p. 128). So although one's awareness of one's auditory experience depends on one's original auditory experience, this original auditory experience could, in principle, occur in the absence of any awareness of it, and so would be an unconscious act of hearing. But Brentano entertains and then systematically rejects a number of the most familiar arguments for the existence of unconscious mental states. So even though he does not rule out the possibility of unconscious mental states *a priori*, he does not think they actually exist. But even if unconscious mental states *do* exist, this does not imperil Brentano's thesis. On the Brentanian model, all *conscious* states have a built-in self-awareness, but one can accept this view while accepting, or remaining agnostic about, the existence of unconscious mental states. To suppose otherwise is, in the apt wording of Thomasson, "to confuse the view that consciousness is intrinsic to those states that possess it, with the different view that consciousness is intrinsic to *all* mental states" (2000, p. 197; cf. Natsoulas, 2004, pp. 102–3). I propose, then, to

24. See also Woodruff Smith (e.g., 1988, 1989), Natsoulas (e.g., 2004), and Kriegel (2003b), all of whom believe that there are unconscious mental states, but defend views that resemble Brentano's.

simply grant, provisionally, that there are unconscious mental states. The Brentanian model remains defensible despite this concession.

By way of enlarging on this claim, let me offer some further support for the claim that every conscious state has an implicit self-awareness. It might be contended that, in this respect, Brentano's account is simply false. There are two versions of this objection.[25] One might insist, in a straightforwardly eliminativist manner, that the kind of ubiquitous non-reflective self-awareness on which Brentano relies simply does not exist, in any form. An ally of this version of the objection contends that self-awareness is a genuine psychological phenomenon, in the sense that we occasionally reflect on our own conscious states, but no implicit self-awareness, either conscious or unconscious, accompanies all our introspective and non-introspective states. This version of the objection amounts to a denial of the Transitivity Principle (see ch. 4), which says that a mental state is conscious only if the subject of the mental state is, in some appropriate way, aware of being in it. On the other hand, one might retain the idea that all conscious states involve self-awareness, but hold that, in being conscious of *x*, the mental state in virtue of which the subject is aware of her consciousness of *x* is itself unconscious. Proponents of higher-orderism fall into this latter category.

Let us address, first, the straight denial or eliminativist version of this objection, according to which implicit self-awareness simply does not exist, conscious or otherwise. Our objector here has no problem countenancing states of consciousness of which the subject is entirely unaware. In response to this version of the objection, I should like to refer back to some results from preceding chapters. In Chapter Two it was argued that conscious experiences are always like something *for* the subject or *to* the subject of the experience. But if a subject is wholly unaware of her experiences, i.e., if she is wholly non-self-aware, there seems to be no sense in which her conscious experiences could be said to be "for her." More rhetorically, how can an experience be something for S or to S if S is not, in some way, aware of that experience? This argument is directly related to the Symmetry Argument, presented in Chapter Four. According to the Symmetry Argument, there is a symmetry between what ordinary objects (events, etc.) are like and what conscious experiences are like, such that if a conscious experience is like something to the subject of the experience, then the subject is aware of that experience. And since every conscious experience is like something to the subject of the experience, it follows that the subject is, in some way and to some degree, aware of the experience.

25. A point also made by Kriegel (2003b, p. 130).

I am also inclined to appeal to authority in response to this version of the objection. Certain contexts non-fallaciously warrant an appeal to authority, or, rather, an appeal to a tradition in which a number of authoritative figures have operated. This is just such a context. A sizeable number of highly influential philosophers and psychologists – including Aristotle (see Caston 2002), Brentano (1874, 1982), Descartes (see Aquila 1988, Woodruff Smith 1993), Fichte (1970), Freud (see Natsoulas 1992), Gurwitsch (1950), Heidegger (see Zahavi 2005), Husserl (see Zahavi 2002b, 2003; Woodruff Smith 2004b), James (1890), Kant (1965), Leibniz (see Gennaro 1999), Locke (1975), Sartre (1956), and, more recently, Armstrong (1968, 1981), Carruthers (1996, 2000), Lycan (1996, 2001), Rosenthal (e.g., 1986, 1990), Woodruff Smith (1986, 1989), and Zahavi (1999, 2004, 2005) – have, in one way or another, subscribed to the Transitivity Principle. What is more, their endorsement of the principle has, for the most part, been unequivocal: they have taken the link between being in a conscious state and being in a self-aware state to be a phenomenological datum, not an expendable theoretical construct (cf. Van Gulick 2004, p. 67). To be clear, the thesis I am attributing to these philosophers is *not* the thesis that all conscious states necessarily involve a *reflective* or *introspective* self-awareness. Rather, I am attributing to them the thesis that a conscious state is always a state of which the subject is, *in some sense*, aware, and this thesis leaves it open whether the subject's awareness of his or her conscious state is a reflective or non-reflective awareness, or even an unconscious awareness. It is possible, of course, that all these clever minds are hopelessly imprisoned within a fictitious paradigm, labouring under some kind of illusion, but that is extremely unlikely.

But what about the second version of this objection, according to which self-awareness is a bona fide feature of conscious states, but the mental states in virtue of which we are aware of our conscious states are themselves unconscious? As we have seen, HO theorists are attracted to – indeed committed to – this idea, arguing that a subject might perfectly well be conscious of x (the target mental state) without the state (the higher-order thought or perception) in virtue of which the subject is conscious of x itself being conscious. But this is highly problematic, in at least three ways. First, as I have already pointed out (ch. 5, §5.2.2), such a view is phenomenologically idle: we are *experientially* aware of our conscious states. Second, suppose that every conscious state is accompanied by, and rendered conscious by, an unconscious state. If this is true, then self-awareness is a theoretical posit, known only by inference from certain observations. We infer, perhaps from behavioural indicators, that we are aware of our conscious states. But this seems manifestly at odds with what self-awareness is and how we acquire knowledge of it. As Woodruff Smith has convincingly argued, the knowledge we have of self-awareness is unmediated, non-inferential knowledge, i.e., *experiential*

knowledge.[26] We *experience* our conscious states, and this experiential evidence grounds our knowledge of the ubiquitous nature of self-awareness.[27] Third, the kinds of cases HO theorists sometimes adduce to bolster the claim that the mental states in virtue of which we are aware of our conscious states are themselves unconscious – e.g., cases like the "cocktail party effect" – cannot play the role that HO theorists assign to them. The cocktail party effect supposedly shows that it is possible for a creature to be conscious of *x* (say the other conversations) without the state in virtue of which the creature is conscious of *x* itself being conscious. However, even if the cocktail party effect shows that the mental states in virtue of which subjects are conscious of *extra-mental objects* can themselves be unconscious, it does not show that the mental states in virtue of which subjects are aware of their own *conscious states* can themselves be unconscious. For if the mental state in virtue of which a subject is aware of his own conscious state is itself unconscious, then it is difficult to see how the subject could make a report on that state. Notice that the subject, in undergoing the cocktail party effect, is unable to say anything about the other conversations prior to hearing his name being uttered. But subjects are usually able to make detailed spontaneous reports on their conscious states, even non-introspected ones, which suggests that the mental states in virtue of which they are aware of those states are themselves conscious.

There are compelling reasons, then, to suppose that the mental states in virtue of which we are aware of our conscious states are themselves conscious, just as Brentano prescribed. Such a view, however, does not contravene the assumption that there unconscious mental states.

Allow me to pause briefly to review progress. Thus far I have explored the principal elements of the Brentanian model of consciousness. It deserves emphasis

---

26. Woodruff Smith writes: "unconscious mental processes and the subtler structures of mental states – the posits of psychoanalysis or cognitive psychology – are indeed theoretical entities, known only by inference from relevant observations. But the basic phenomena of mind and self, propositional attitudes and qualia, consciousness and self-consciousness – these things we know from our own experience, from our inner awareness of them in consciousness" (1986, p. 154).

27. It should be noted that Woodruff Smith no longer holds the view that all conscious states have this kind of phenomenologically revealed self-awareness. In his recent work he argues that we "experience very basic modes of consciousness that lack 'self-consciousness' " (Woodruff Smith 2004b, p. 109). Examples of such states of consciousness include hammering a nail or choosing to hit a tennis ball crosscourt rather than down the line. In these cases, Woodruff Smith argues, even though I am fully conscious, I am aware neither of myself nor of my mental activity. I disagree. Even when one is hammering a nail, one's experience is phenomenologically revealed, in the sense that it is "manifested" or "given" to the subject in a *non-objectual* way. I argue for this claim below (§6.2.2).

that, as my defence of it suggests, I lean toward believing that Brentano's account is highly plausible. There is a powerful argument in its favour, and it seems to have the resources to avoid those difficulties that are endemic to the HO theory. Since Brentano does not posit two distinct levels of mental activity, there is no threat of confabulated conscious states. In addition, the Brentanian model has the desirable property of being phenomenologically sensitive: it respects the fact that the mental states in virtue of which we are aware of our conscious states are themselves conscious.

6.2.2 A twist

At this point it is necessary to anticipate and confront a response that my discussion of the Brentanian model might provoke. The response involves the counterclaim that Brentano's account of state consciousness ought to be rejected, since it makes the fatal mistake of construing implicit self-awareness in terms of a "subject-object" relation, i.e., in terms of conscious states being objects to which subjects stand in a certain kind of relation. Zahavi has objected vigorously to the Brentanian model on these grounds (e.g., 1999, 2004, 2005),[28] arguing that a more promising alternative to the HOT theory can be found in the work of, among others, Husserl and Sartre, who do not make this mistake. According to Sartre, for example, even though our conscious states themselves are always "given" or "manifested" in experience, they do not become *objects* of consciousness except during acts of reflection: "this consciousness of consciousness – except in the case of reflective consciousness – is not *positional*, which is to say that consciousness is not for itself its own object" (Sartre 1957, pp. 40–1). Thus, Sartre does not deny that consciousness involves self-awareness, but he denies that self-awareness can be accounted for on analogy with our consciousness of extra-mental objects, i.e., in terms of a subject-object relation.

There is, it seems to me, something fundamentally right about Sartre's position: our conscious states simply aren't presented to us in the same way extra-mental objects are. Any consciousness of the subject-object variety "operates," to use Zahavi's words, "with a duality" or an "epistemic *difference* between the subject and the object of experience" (2004, p. 70). But when I consciously perceive *x*, there is no distance, as it were, between my perceiving *x* and my being implicitly aware of perceiving *x*. To be sure, we can, as Sartre would put it (e.g., 1956, p. liii), take up an explicit or reflective position vis-à-vis our conscious states and *posit* them as objects, in which case they become objects of consciousness. But implicit self-awareness is not a kind

28. See also Thomasson 2000, Thomas 2003, and Drummond 2006.

of reflection. Rather, to the extent that a mental state is conscious it is *lived through*, and so the subject is immediately *acquainted* with it, though not as a kind of object (cf. Chalmers 1996, pp. 196–7). Zahavi has employed a useful metaphor to illustrate this point. "Consciousness," he says, "can. . .be compared to a flame, which illuminates other things, and itself as well" (1999, p. 34).[29]

It is most natural and intuitive to think of our perceptual and non-perceptual experiences as constituting the means *through which* we are conscious of things. Experiences are not objects – for example, my visual experience of *x* is not an item of which I am aware in being visually conscious of *x*. Rather, my experience enables me to gain visual access to *x*: *x* is present to me, *through* my seeing it (cf. Zahavi 2005, pp. 121-2). But my experience is still "there" all along; it is manifested to me. And I may, if I wish, turn my attention toward it and reflect on it.

Notice that it will not do – as, for example, Kriegel (e.g., 2003a,b,c) does – to gloss implicit self-awareness in terms of the distribution of the subject's attention resources, i.e., in terms of the distinction between *focal* and *peripheral* awareness.[30] This is because focal and peripheral awareness are both kinds of *object-awareness*, but my implicit awareness of my visual experience is not akin to the kind of inattentive visual awareness I have of the printer beside my computer. As Woodruff Smith has put the point: "my immediate awareness of my experience is not a second-order observation, not even one with a lower degree of attention" (1989, p. 87).

Apart from its phenomenological attractiveness, there is a simple yet compelling argument for the claim that implicit self-awareness cannot be construed along subject-object lines. To wit, if implicit self-awareness is a kind of object-consciousness,

---

29. McGinn also seems sympathetic to the view that subjects are acquainted with their experiences in this way. He argues that the "subjective aspects of experience involve reference to the subject undergoing the experience – this is what their *subject*ivity consists in. But we can also say that perceptual experiences have a world-directed aspect: they present the world in a certain way, say as containing a scarlet sphere against a blue background. This is their representational content, what states of affairs they are *as of*. Thus perceptual experiences are Janus-faced: they point outward to the external world but they also present a subjective face to their subject; they are of something other than the subject and they are like something for the subject" (1991b, p. 29).

30. Curiously, Kriegel claims that, in characterizing implicit self-awareness, the distinction between focal and peripheral awareness is a distinction that "virtually all phenomenologists interested in self-awareness appeal to" (2003c, p. 176). But he is simply wrong about this. Certainly Husserl, Heidegger, Sartre, and Merleau-Ponty, four of the prime figures of twentieth-century phenomenology, and all keenly interested in self-awareness, never appealed to the distinction. Indeed, for reasons already provided, Husserl (at least in his later writings) and Sartre would reject it outright as a way to characterize implicit self-awareness. What is more, some phenomenologists (e.g., Woodruff Smith 1989, p. 87) writing prior to Kriegel's invocation of the distinction expressly argued *against* characterizing implicit self-awareness in such terms.

then conscious states are mental objects. However, since conscious states are not mental objects, implicit self-awareness is not a kind of object-consciousness. Many philosophers are unfavourably disposed to the key premise in this argument – i.e., that conscious states are not mental objects – but I would surmise that this is due, at least in part, to the misleading surface structure of the grammar of perception sentences. Expressions like "I had a perceptual experience" and "I am having a perceptual experience," i.e., expressions in which the term "perceptual experience" is grammatically the accusative object, are perfectly licit. But we should not be misled by the grammar of these expressions. To *have a perceptual experience* is simply to be in a conscious perceptual state; it is not to possess an object that happens to be a perceptual experience (unlike, say, to *have a wristwatch* is to possess an object that happens to be a wristwatch).[31] And to be in a conscious perceptual state is simply to be perceptually conscious of something. We can, to repeat, reflect on our perceptual experiences, in which case the experience becomes the object of one's act of reflection. But our ability to reflect on our perceptual experiences is no reason to suppose that, prior to reflection, they are mental objects.

Now Brentano frequently speaks of primary and secondary *objects* of awareness, and so he is reasonably interpreted as being mired in the object-consciousness paradigm, but we can abstract away from this potentially contentious aspect of his position without abandoning his core insight, and that is exactly what I propose to do. Notice that, as I have expounded it, the core insight of the Brentanian account is that there is a single token state involved in one's perceiving *x* and one's being implicitly aware of perceiving *x* (cf. Thomasson 2000, pp. 193–4). But Sartre would agree with this point. He repeatedly emphasizes that the relation between one's consciousness of *x* and one's implicit awareness of one's consciousness of *x* is non-dyadic. We are told, for example, that "To be and to be aware of itself are one and the same thing for consciousness" (1957, p. 83), and that "We understand now why the first consciousness of consciousness is not positional; it is because it is one with the consciousness of which it is consciousness" (1956, p. lvi). For Sartre, then, just like Brentano, the mental state M in virtue of which the subject of M is aware of M is M itself; only, for Sartre, it is wrong to construe the self-awareness involved in M as a peculiar kind of object-consciousness.

One might object that it is inconsistent to claim that we are aware *of* our experiences and that our experiences are not given as objects. But Sartre replies that the "necessity of syntax has compelled us hitherto to speak of the 'non-positional consciousness of self'" (1956, p. liv). Thus, Sartre explicitly acknowledges that his

31. Cf. Ryle 1949, esp. pp. 208–9; Bennett and Hacker 2003, esp. §3.8; and Douglas 1998.

use of "of" is illicit, since it suggests that we are aware of our experiences in the same way we are aware of extra-mental objects, but his (and my) usage of the term need not cause confusion.[32]

I conclude that the core insights of the Brentanian and Sartrean models can be combined to yield a well-motivated account of what constitutes a conscious state.

Some dense philosophical thickets lie ahead, however. On the model of consciousness I have been discussing, when one is conscious of *x*, one is conscious not only of *x*, but also, to some degree, of one's consciousness of *x*; that is, our conscious states themselves are always pre-reflectively and non-objectually given or manifested in experience. There is reason to suppose, however, that conscious states are more complex than this description suggests, i.e., that they possess what might be called *de se* content. Roughly characterized, *de se* content is first-person content or "content that concerns not the world but oneself" (Lewis 1999, p. 317).[33] Both Brentano and Sartre, at least in their mature philosophies,[34] held that conscious experiences have a *de se* component. That is, they held that my implicit awareness of a particular experience includes, *ipso facto*, an awareness of myself as subject of that experience. On this view, one is not merely implicitly aware of *seeing* the tree; rather, one is implicitly aware of *oneself seeing* the tree. The inclusion of *de se* content adds a complex, and controversial, dimension to the Brentanian/Sartrean model of consciousness I have sketched.[35]

6.2.3 *De se* content

Not all supporters of the view that implicit self-awareness is an intrinsic feature of conscious states share the belief that it is necessary to posit *de se* content. According

---

32. Sartre (1956, p. liv) distinguishes implicit self-awareness from the reflective self-awareness of phenomenology by referring to the former as "*conscience (de) soi*" ("de" in brackets) and the latter "*conscience de soi*" ("de" not in brackets). The brackets are meant to indicate that implicit self-awareness is not a positional consciousness relating a subject to an object.

33. For further critical discussion of *de se* content, see Castañeda 1966, Chisholm 1981, Gennaro 1996, McGinn 1983, and Perry 1979.

34. See Brentano 1874, pp. 276–7, and Sartre 1956, pp. 162–3.

35. Some HOT theorists have also argued that conscious states have *de se* content. According to Rosenthal, for example, "[a] mental state's being conscious consists of being conscious of oneself as being in that state. A state is not a conscious state if I am just conscious of it, on its own, and not conscious of myself as being in the state" (2002b, pp. 242–3; see also Gennaro 1996, ch. 8). However, since higher-orderism has been shown to be defective on grounds that are quite independent of any *de se* considerations, I propose to forgo discussion of how it fares with respect to this aspect of consciousness.

to Kapitan, for example, "The fundamental reflexive encounter that accompanies my conscious states is not *first-person-awareness*; it is awareness of the conscious activity alone, without any identification as being me, mine, or anything else" (1999, p. 40). But I cannot see how this view can be defended. In hearing something (seeing something, etc.), one is, as I have argued, aware of one's act of hearing. But an act of hearing is (whatever else it is) an act of a subject, i.e., an act of a subject toward an object. But if so, and if one's implicit awareness of one's act of hearing is an intrinsic component of that act of hearing, then it is difficult to see how one could be implicitly aware of that act of hearing without being implicitly aware of oneself hearing.

Moreover, invoking *de se* content seems necessary in order to explain certain behaviours and behavioural dispositions in certain contexts.[36] For example, if a dog sees a cat up in a tree, the dog's visual state must make it possible for the dog to relate *itself* with respect to the cat; otherwise we could not explain the dog's readiness to pounce on the cat. Presumably, then, the dog is not merely implicitly aware of *a* visual experience of the cat in the tree, but also implicitly aware of *itself* standing in a unique spatial relation to the cat in the tree. This *de se* content must be part of the dog's visual consciousness because it helps explain the dog's behavioural disposition.

We can take a step toward making the *de se* aspect of state consciousness more precise by calling attention to the fact that some commentators construe this aspect of it in terms of *ownership*. The idea is that all conscious states are "owned" and are experienced as owned: in consciously seeing *x*, I am implicitly aware of *my* seeing *x*. Woodruff Smith provides the following apt description of this idea:

> . . . as I see that dove, I am immediately aware of *my* seeing that dove: I am aware that "I see that dove". My awareness of the experience does not leave it open whose experience this is; it is mine. If you will, the intentional character of the experience is "directional", aiming as if *from* a subject *toward* an object. And my inner awareness of the experience reveals that directional structure, from "me" or "I" toward "this dove".
>
> (1989, p. 88)

And Kriegel says:

> When we have a mental state self-consciously, there is a subtle awareness of self implicit in that state, whereby we are aware of ourselves as its *owners*. This is not just a matter of the traditional issue of ownership; it is a matter of ownership-awareness. It is not just that we *are* the owners; we are *aware* of being the owners, albeit implicitly and inattentively.
>
> (2003b, p. 104)

---

36. See Gennaro 1996, pp. 172–3.

Now, interestingly – and, as I will explain below, not unproblematically – both Woodruff Smith and Kriegel take implicit *ownership* awareness to be a simple means of elucidating the claim that, in being conscious of *x*, one is implicitly aware of *oneself* as being conscious of *x*. Woodruff Smith is unambiguous about this:

> Inner awareness of one's passing experience is ipso facto awareness of oneself as subject of that experience. . . . I am not merely aware of a passing experience, this visual appearing of that dove, without reference to any subject. Rather, as I see that dove, I am immediately aware of *my* seeing that dove.
>
> (1989, p. 88)

So is Kriegel:

> In your auditory experience of the bagpipe you are aware primarily, or *explicitly*, of the bagpipe sound; but you are also *implicitly* aware that this auditory experience of the bagpipe is *your* experience. That is, you are aware of yourself as the subject of [the] experience.
>
> (2003b, p. 104)

Both Woodruff Smith and Kriegel, then, equate implicit awareness that the experience of *x* is *yours* with implicit awareness of *oneself* as the subject of the experience.[37] But this seems to gloss over a crucial distinction, for surely there is a (non-trivial) difference between the following two positions: (A) in consciously seeing *x*, I am implicitly aware of *myself* seeing *x*, and (A′) in consciously seeing *x*, I am implicitly aware of the seeing as *my* seeing *x*. These positions are not equivalent because (A) need not involve any awareness of ownership, at least if we are using the term "ownership" in its ordinary sense, whereas (A′), of course, does.

In fact, there are a number of different positions in the offing here. Consider, first, the following positions, which do *not* involve *de se* content:

(DS_1) A subject S consciously perceives *x* at *t* only if at *t* S is implicitly aware of perceiving *x* (i.e., a perceptual state P in virtue of which a subject S is conscious of *x* at *t* is conscious only if at *t* S is implicitly aware of perceiving *x*).[38]

(DS_2) A subject S consciously perceives *x* at *t* only if at *t* S is implicitly aware of a perception of *x*.

---

37. See also Zahavi 2005, pp. 124–5.

38. I add this parenthetical clause to forestall confusion. For ease of exposition I will use the language of transitive creature consciousness to describe the various positions involving, or not involving, *de se* content, but this should not be construed as an attempt to articulate the conditions under which we are justified in affirming that a creature is conscious of something. My concern here, as always, is with conscious *states*, this time with their *de se* content.

It is not clear whether DS_2 is not just an alternative rendering of DS_1. In DS_1 the subject is implicitly aware of an activity, i.e., an activity of perceiving *x*, whereas in DS_2 the subject is implicitly aware of an "object," namely, a perception of *x*. But a case could be made that there is no difference between a perception of *x* and an activity of perceiving *x*. At any rate, neither of these positions involves *de se* content, because the subject of the experience is not implicitly aware of *herself* perceiving *x*, or of *her* perception of *x*, but rather simply implicitly aware of *an* activity of perceiving *x*, or of *a* perception of *x*. However, as I have argued, the having of some kind of *de se* attitude is presupposed in the having of a conscious experience, so these positions, since they ignore *de se* content, are inadequate.

The introduction of *de se* content complicates matters considerably. We can distinguish the following positions:

List A

(DS_3) A subject S consciously perceives *x* (has a conscious sensation, etc.) at *t* only if at *t* S is implicitly aware of *herself*.

(DS_4) A subject S consciously perceives *x* at *t* only if at *t* S is implicitly aware of *herself perceiving x*.

(DS_5) A subject S consciously perceives *x* at *t* only if at *t* S is implicitly aware of herself *qua* perceiving *x*.

(DS_6) A subject S consciously perceives *x* at *t* only if at *t* S is implicitly aware of the perceiving as *her* perceiving *x*.

(DS_7) A subject S consciously perceives *x* at *t* only if at *t* S is implicitly aware *that* she is perceiving *x*.

(DS_8) A subject S consciously perceives *x* at *t* only if at *t* S is implicitly aware *that* the perceiving is *her* perceiving *x*.

Notice that the conscious states at issue in this list of positions are non-propositional; that is, they are conscious states, such as simple seeings, hearings, sensations, etc., that do not, arguably, involve the deployment of concepts (though the conscious states in DS_7 and DS_8 have an implicit self-awareness that is propositional, and since implicit self-awareness is an intrinsic feature of conscious states, those conscious states are also partly propositional). However, if I consciously perceive *that x* – e.g., if I consciously see that the grass is green, or if I consciously think that Russell's paradox is insoluble – then I must have the ability to wield whatever concepts are involved in the sentence following the "that" clause, that is, I must have the ability to wield whatever concepts are involved in *x*. Thus, List A forms only a subset of the available positions, for there are a number of positions in the

offing regarding conscious states that have propositional content. Consider the following list:

List B

(DS_9) A subject S consciously perceives that *x* (consciously thinks that *x*, etc.) at *t* only if at *t* S is implicitly aware of *herself*.

(DS_{10}) A subject S consciously perceives that *x* at *t* only if at *t* S is implicitly aware of *herself perceiving x*.

(DS_{11}) A subject S consciously perceives that *x* at *t* only if at *t* S is implicitly aware of herself *qua* perceiving *x*.

(DS_{12}) A subject S consciously perceives that *x* at *t* only if at *t* S is implicitly aware of the perception of *x* as *her* perception of *x*.

(DS_{13}) A subject S consciously perceives that *x* at *t* only if at *t* S is implicitly aware *that* she is perceiving *x*.

(DS_{14}) A subject S consciously perceives that *x* at *t* only if at *t* S is implicitly aware *that* the perception of *x* is *her* perception of *x*.

I will offer arguments indicating that the view of implicit self-awareness captured in DS_4 from List A is the most plausible position to adopt with respect to conscious states that lack propositional content, whereas the view of implicit self-awareness captured in DS_{11} and DS_{13} from List B are the most plausible positions to adopt with respect to conscious states that have propositional content. I shall commence by eliminating some of the candidate positions from each list. A number of positions stand and fall together; for example, position DS_3 from List A is, I will argue, objectionable, but the consideration that undermines DS_3 also undermines its propositional counterpart – i.e., position DS_9 – from List B. After eliminating some of the positions from each list, I will attempt to show that the remaining positions (with one exception) are highly attractive, not only because they can dodge the various difficulties attending the other positions, but also because they are consonant with an approach to the content of implicit self-awareness, developed on the Brentanian/Sartrean model, that is itself plausible and independently motivated.

Let us begin, then, with position DS_3 from List A. According to this position, if I consciously perceive *x*, then I am implicitly aware of myself. But this is quite implausible, for if, in consciously seeing *x*, I am merely implicitly aware of myself, then my conscious state, i.e., my state of seeing *x*, is not a conscious one, since it is not a state of which I am aware. Position DS_9 from List B fails for the same reason.

Position DS_6 from List A is equally faulty. On the one hand, if, as Kriegel and Woodruff Smith seem to assume, saying that I am implicitly aware of *my* seeing *x* is

just an alternative way of saying that I am implicitly aware of *myself* seeing *x*,[39] then DS_6 is defensible, since, in that case, it is equivalent to DS_4. On the other hand, if DS_6 is construed as the claim that, in consciously seeing *x*, I am implicitly aware of my state of seeing as being a perceptual state that I *possess*, then the position makes scant sense. One cannot possess a conscious experience in the same way that one can possess an automobile or a toothbrush. Conscious states are things one can be said to be *in*, and conscious experiences are things one can be said to *have*, in the sense of *undergo*. But it makes no sense to say that one owns one's conscious experiences.[40] Thus, unless "ownership" is being used in a (as yet undefined) special technical sense, it seems false to say that we have an implicit awareness of experience ownership. Position DS_6, then, from List A, along with position DS_8 from List A and positions DS_{12} and DS_{14} from List B, are rejectable, since each of these

---

39. Zahavi endorses unmistakably the view that saying I am implicitly aware of *my* seeing *x* is just an alternative way of saying that I am implicitly aware of *myself* seeing *x*. He holds that all conscious experiences have an implicit quality of *mineness*: "When I. . . am aware of an occurrent pain, perception, or thought from the first-person perspective, the experience in question is given immediately, noninferentially and noncritically as *mine*" (2005, p. 124). However, he is quick to emphasize that a conscious experience is not something one *possesses*: "Experiences are not something one simply has, like coins in the pocket" (2005, p. 116). For Zahavi, then, it is not as though, in having a conscious experience, one is tacitly aware of experience *ownership*. On the contrary, he takes tacit awareness of mineness to be a means of clarifying the claim that, in being conscious of *x*, one is tacitly aware of *oneself* as being conscious of *x*. Indeed, he *identifies* this "pre-reflective sense of mineness with a minimal, or core, sense of self" (2005, p. 125).

40. Consider, in this connection, Wittgenstein's remark on the following questions:

(1) "Are these books *my* books?"
(2) "Is this foot *my* foot?"
(3) "Is this body *my* body?"
(4) "Is this sensation *my* sensation?"

In regards to (4), Wittgenstein says: "Which sensation does one mean by '*this*' one? That is: how is one using the demonstrative pronoun here? Certainly otherwise than in, say, the first example!" (1953, §411). And cf. Bennett and Hacker: "Having a pain is not owning or possessing anything – any more than having a train to catch is owning or possessing something. Having a pain is *not* standing in any special relation to a pain, since it is not a relation at all – any more than is having a depression, which is no more than a matter of feeling depressed. To have a pain in one's foot is a matter of one's foot hurting" (2003, p. 96). They go on to say: "We are inclined to think otherwise, because we unwittingly construe *having a pain* on the model of *having a penny*. But this is mistaken, for *having a penny* is indeed a case of ownership or possession, and for a person to have a penny is for him to stand in a relationship – namely, of ownership or possession – to the penny. But for a person to have a pain is not for him to stand in a relation to a pain" (2003, p. 96). I concur wholeheartedly with Wittgenstein and Bennett and Hacker on these points.

positions relies on the implausible assumption that the *de se* content of implicit self-awareness involves some kind of awareness of experience ownership.

DS_7 from List A is also vulnerable to serious criticism. This position suggests that all conscious experiences, including relatively simple perceptual experiences, have a *de se* component that is cognitive or propositional, such that in seeing *x* I am implicitly aware *that* I am seeing *x*. But the intellectual requirements thereby made requisite for consciousness would seem to preclude the existence of consciousness in infants and animals.[41] In order for me to be implicitly aware *that* I am seeing *x*, I require the concept of seeing, as well as any concepts involved in *x*. But it is safe to assume that non-discursive creatures are normally incapable of employing such concepts. A mouse, for example, is capable of consciously smelling cheese, but presumably the mouse is incapable of being implicitly aware *that* it is smelling cheese. DS_7, then, can be jettisoned, since it compels one to deny, unreasonably, that infants and most animals have conscious experiences.

The reality of infant and animal consciousness also compromises the viability of position DS_5 from List A. According to this position, if I consciously perceive *x*, then I am implicitly aware of myself *qua* perceiving *x*. However, being implicitly aware of myself qua perceiving *x* would seem to require the same level of intellectual wherewithal that position DS_7 requires. For example, in being implicitly aware of myself qua seeing *x*, I must have the intellectual ability to subsume my visual state under some category – the "visual state" category, say – which would involve having the intellectual ability to distinguish my visual state from my other perceptual states. But, again, it is doubtful that infants and animals possess such sophisticated intellectual resources. To be sure, infants and animals see, hear, smell, etc. a whole host of objects, features, and states of affairs, but presumably they are not implicitly aware of their perceptual states *qua* being states of a specific kind.

The considerations advanced so far indicate that positions DS_3, DS_5, DS_6, DS_7, and DS_8 from List A and positions DS_9, DS_{12}, and DS_{14} from List B are

41. Some philosophers (e.g., McDowell 1994) hold that even the most simple perceptual states are laden with concepts of some sort, even if merely elementary recognitional and discriminatory concepts. These philosophers identify concept possession in its simplest forms with bare recognitional capacities, and accordingly attribute concept possession to creatures that exhibit behaviourally the right sort of recognitional capacities. I am more inclined, myself, to opt for a linguistic construal of concept possession, according to which, in Bennett and Hacker's words, "*To have a concept is to have mastered the use of a word (or phrase)*" (2003, p. 340). It is sometimes claimed that this construal exaggerates the importance of language, but surely not *all* organisms and artefacts, such as garden slugs and vending machines, that respond discriminatively to classes of objects possess concepts, and I cannot see how those who endorse the more broad-spectrum construal can exclude granting concepts to those things.

unsatisfactory. We are left, then, with position DS_4 from List A and positions DS_{10}, DS_{11}, and DS_{13} from List B. As I said, these positions are worth examining more carefully not only because they can deflect the various charges of inadequacy that can be levelled against the other positions, but also because a particular account of the content of implicit self-awareness, developed on the Brentanian/Sartrean model, strongly supports the claim that these positions, with the exception of DS_{10} from List B, are probably true. I will first put forward the account of the content of implicit self-awareness developed on the Brentanian/Sartrean model. To be succinct, I will refer to it as MB (for "modified Brentanian account of the content of implicit self-awareness"). After expounding MB, I will show how such an account, itself well-motivated, supports positions DS_4, DS_{11}, and DS_{13}.

It seems to be an unreflective prejudice among higher-order theorists that the content of higher-order awareness must be either cognitive or experiential,[42] but not both. This prejudice explains why higher-orderism has been rent by a kind of schism between those who believe that the content of higher-order awareness is purely experiential (e.g., HOP theorists such as Armstrong and Lycan) and those who believe it is purely cognitive (e.g., HOT theorists such Rosenthal, Gennaro, and Carruthers). If one accepts MB, however, disputing the merits and demerits of these respective positions is unnecessary, since, according to MB, the content of implicit self-awareness can be either experiential *or* cognitive *or* both, depending on whether the conscious state in question is experiential or cognitive or both.

Let me explain. Recall that, according to Brentano, the secondary content of any conscious state *c* is *c* itself. This is the case for simple perceptual states as well as for more cognitively demanding propositional states:

> No one who recalls our discussion of presentations will maintain that assuming that each mental phenomenon is the object of an accompanying cognition leads to an infinite complication of mental states and therefore would be by its very nature impossible. The characteristic fusion of consciousness and the object of consciousness is just as evident in cognition as it was there. Whenever a mental act is the object of an accompanying inner cognition, it contains itself in its entirety as presented and known, in addition to its reference to a primary object.
> (1874, pp. 138–9)

But this view has implications for how we view the content of implicit self-awareness, for, when combined with the Sartrean non-object account of experiential self-givenness, it entails that if one's experience is cognitive, then one's implicit

42. I here use "cognitive" and "experiential" as umbrella terms for any expressions used to denote, respectively, states that require the deployment of concepts and states that do not.

awareness of one's experience will likewise be (at least partly) cognitive, since the experience – a cognitive one – is immediately and non-objectually given to the subject. For the same reason, if one's conscious state is experiential, then one's implicit awareness of one's experience will likewise be (at least partly) experiential. For example, if I consciously see a blue sky, I "live through" my visual experience such that my visual experience itself is immediately and non-objectually manifested in my experience. In this case, my implicit awareness of my conscious state is at least partly experiential, since my conscious state itself is experiential. In contrast, if I see *that* the sky is blue, my implicit awareness of my conscious state is at least partly cognitive, since my conscious state itself is partly cognitive. In short, my implicit awareness of a particular conscious state S is derived from S itself, which means that S *just is* the content of my implicit self-awareness. It follows that the content of the implicit self-awareness involved in conscious states can be either experiential or cognitive.

Now Brentano himself held that one's implicit awareness of oneself seeing, hearing, smelling, etc. is both a form of perception *and* a form of judgement or thought, so that *all* conscious states are directed upon some primary object, but are at the same time experientially (or, in Brentano's jargon, presentationally) and cognitively (or, in Brentano's jargon, judgementally) directed upon themselves (see Brentano 1874, bk. II, ch. III, §1-3).[43] Brentano posits this multipart conception of implicit self-awareness primarily because he believes that consciousness is a form of knowledge:

> What is certain is that knowledge frequently accompanies mental phenomena. We think, we desire something, and know that we think and desire. But we only have knowledge when we make judgements. It is beyond doubt, therefore, that in many cases along with the mental act there exists within us only a presentation which refers to it but also a judgement about it.
>
> (1874, p. 138)

Quite apart from its phenomenological dubiousness, however, such a conception of implicit self-awareness seriously over-intellectualizes the implicit self-awareness required for consciousness. If it is true, very few creatures would qualify as subjects of

43. In fact, Brentano originally held that all conscious states have a fourth component, viz., a feeling of pleasure or displeasure (1874, bk. II, ch. III, §5–6). Thus, he held that all conscious states are primarily directed at some intentional object, but are also experientially, cognitively, and pleasurably or un-pleasurably directed upon themselves. Wisely, he later abandoned this fourfold view of conscious states (see Brentano 1874, p. 276), though he continued to adhere to the threefold view, i.e., that conscious states are primarily directed upon some intentional object, but are also experientially and cognitively directed upon themselves.

conscious experience. It is safe to assume, for example, that a cat, in seeing a mouse, does not make an implicit judgement that it is seeing a mouse.

Brentano's response to this difficulty consists in denying that the judgements involved in implicit self-awareness are conceptually very sophisticated:

> . . .this judgement of inner perception does not consist in the connection of a mental act as subject with existence as predicate, but consists rather in the simple affirmation of the mental phenomenon which is present in inner consciousness.
>
> (1874, p. 142)

So, according to Brentano, a subject, in seeing something (hearing something, etc.), makes an implicit judgement that consists simply in affirming that his act of seeing is occurring.

Even this sort of judgement, however, requires possession of concepts that it would be implausible to ascribe to most species of animal. After all, to implicitly affirm that one's act of seeing is occurring, one requires the concept of seeing, but dogs, cats, birds, etc. do not possess the concept of seeing. Now it seems that, on Brentano's view, one does not have to be capable of entertaining a judgement about one's state of seeing *qua* a state of seeing. He seems to suggest in the quotation cited above that it might be sufficient if one were capable of thinking of one's state of seeing *as distinct* from some other state. Thus, the relevant implicit judgement might take the form: <I am having an experience, and *this* experience is distinct from *that* experience>. But even here the subject needs to employ a relatively sophisticated concept of experience, for if the subject distinguishes between "*this*" and "*that*" judgementally (i.e., cognitively), then whatever "*this*" is must be distinguishable from whatever "*that*" is by means of a concept that distinguishes "*this*" from "*that*." But "*this*" is a type of experience, and so is "*that*," which means that the concept that distinguishes "*this*" from "*that*" would have to be a concept of experience.[44] On the other hand, if one denies that a concept of experience is required to make the relevant distinction, then it would seem that the kind of distinction being made is not one that requires any conceptual apparatus after all; rather, the distinction being made is one that is purely experiential. But if a creature is capable of implicitly distinguishing the kinds of conscious states it has purely experientially, then there is no need to posit the existence of an implicit judgement that accompanies every conscious state in the first place.

I think it best to by-pass all danger of attributing exceedingly high-level conceptual powers to infants and animals. We should therefore avoid Brentano's claim

---

44. This is a version of an argument found in Carruthers (2000, p. 195), though Carruthers employs the argument for quite different purposes.

that all conscious states involve both an experiential and cognitive self-awareness. To be sure, some conscious states – for example, concept-imbued perceptual states – have a content that is both experiential and cognitive, which means that the content of implicit self-awareness can be both experiential and cognitive. But, for reasons given above, it seems incorrect to claim, as Brentano does, that the content of the implicit self-awareness of *every* conscious state is both experiential and cognitive.

According to the modified Brentanian account of the content of implicit self-awareness (MB) I want to propose, then, our conscious states themselves determine the content of implicit self-awareness. The account may be explicated as follows. If I make a non-perceptual judgement, if, for example, I judge that Russell's paradox is insoluble, then the implicit self-awareness involved in my judgement will be purely cognitive, since such a non-perceptual judgement is purely cognitive. On the other hand, if I see (hear, etc.) that something is the case, then the implicit self-awareness involved in my perceptual state will be both experiential and cognitive, since seeing *that* something is the case is a concept-imbued perceptual state – a perceptual judgment, if you will. In this case, the cognitive content issues from the judgement, whereas the experiential content issues from the seeing. Finally, if I simply see an object on which I bring no concept to bear, then the implicit self-awareness involved in my state of seeing will be purely experiential, since simple seeings, hearings, etc. do not require the mobilization of concepts. On this account, the content of implicit self-awareness can be experiential or cognitive or both, depending on the conscious state in question.

So let us take stock. A *prima facie* consideration in favour of MB is that it allows us to adjudicate the controversy over whether the content of one's implicit self-awareness is experiential or cognitive. In addition, it affords several advantages over Brentano's original account. First, according to MB, the content of implicit self-awareness is less complex than on Brentano's account. Someone sympathetic to Brentano's account is entitled to rejoin, of course, that Brentano's account is only as complex *as is necessary*, but the arguments I have examined indicate that the content of implicit self-awareness is not as complex as Brentano's account suggests; and surely it is wiser to opt for the simpler alternative if there is nothing lost in doing so. Second, MB is more phenomenologically sensitive: phenomenologically speaking, we do not make implicit conceptual judgements every time we perceive something, have a sensation, etc. Third, unlike Brentano's original account, MB avoids any difficulties that might arise in connection with having to attribute sophisticated conceptual powers to infants and brutes.

So where does this leave us with respect to the remaining position from List A – i.e., DS_4 – and the remaining positions from List B – i.e., DS_{10}, DS_{11}, and DS_{13}? As I mentioned, I want to claim that DS_4, DS_{11}, and DS_{13} are plausible, whereas DS_{10} is not. A simplified version of my argument for this claim is that since MB is

well-motivated, it is reasonable to suppose that a view of implicit self-awareness that satisfies its requirements is itself well-motivated; and, since DS_4, DS_{11}, and DS_{13} satisfy its requirements, whereas DS_{10} does not, DS_4, DS_{11}, and DS_{13} are well-motivated, whereas DS_{10} is not.

Defending this argument is basically a matter of showing how, unlike position DS_{10}, positions DS_4, DS_{11}, and DS_{13} satisfy MB's requirements. I begin with DS_4, which says that, in consciously perceiving *x*, I am implicitly aware of myself perceiving *x*. This position accords well with MB, since, according to MB, if the primary content of a conscious state is wholly non-cognitive, then the implicit self-awareness involved in that state is wholly non-cognitive; and DS_4 says that one's wholly non-cognitive perceptual state is accompanied by a wholly non-cognitive implicit self-awareness that has a built-in *de se* attitude.

Turning to DS_{11} and DS_{13}, these positions are very close neighbours; in fact, I am inclined to think that distinguishing them would be somewhat pedantic, since being implicitly aware *that* I am perceiving *x* and being implicitly aware of myself *as* perceiving *x* are very similar implicit perceptual judgements. In any case, both positions are attractive, since, in accordance with MB, they state that if the primary content of one's conscious state is both experiential and cognitive, then the content of one's implicit awareness of that state is likewise both experiential and cognitive.

On the other hand, position DS_{10} – according to which, in consciously perceiving that *x*, I am implicitly aware of *myself perceiving x* – is incompatible with MB. If MB is correct, then a cognitive state, such as the one described in DS_{10}, would involve an implicit self-awareness that is correspondingly cognitive. But the implicit self-awareness involved in DS_{10} is non-cognitive: to be implicitly aware of oneself perceiving *x*, does not require the employment of concepts.

To summarize, then, MB is a plausible account of the content of implicit self-awareness, and since the views of implicit self-awareness captured in DS_4, DS_{11} and DS_{13} satisfy MB's requirements, DS_4, DS_{11} and DS_{13} are themselves plausible.

6.2.4 The "self"

I want now to make some very brief remarks on the notion of "self" at issue in the preceding discussion, if only to guard against a possible misconception. Even though some sense of "self," according to MB, is automatically implicated in all conscious states, it does not follow that MB is committed to the existence of an underlying, enduring "self" in anything like the Cartesian sense. Nor does it follow that MB is committed to attributing an implausibly strong type of self-awareness to brutes, i.e., one that involves recognizing oneself *as a self*, which no doubt requires the possession of fairly sophisticated concepts. In fact, MB need not posit

a conception of self any more sophisticated than a conception of self that allows a creature to distinguish between itself and other things. In being conscious of *x*, a subject is implicitly aware of herself as being in a particular state of consciousness and this weak form of self-awareness need not involve anything more than knowing the difference between one's body and other things. There are, presumably, certain creatures that live almost entirely in the present, occupying brief states of consciousness that are quickly forgotten. These creatures possess a very minimal conception of self, one that allows them, as physical things, to distinguish between themselves and other objects in the physical world. And this minimal sense of "self" is all the modified Brentanian/Sartrean approach to consciousness that I have been defending requires.

Still, the claim that even infants and brutes are implicitly aware of themselves as being conscious of things is sure to strike some as too strong, for it seems possible to imagine conscious states – those of very young infants, for example – that do not involve the presence of even a minimal sense of self. This is a difficult issue, mainly because it is not clear how much of it can be settled *a priori*. To avoid "Ninth Planet Hegelianism,"[45] therefore, I shall not attempt to make any definitive pronouncements about when children first begin to distinguish between "self" and "other." Nonetheless, at least one *a priori* consideration that may be pressed into service here is that even the unarticulated, "vague" conscious experiences of an infant are not, so to speak, "subjectless"; indeed, conscious experiences cannot *be* subjectless, for conscious experiences are first-personal states. It seems plausible, therefore, that even an infant's conscious experiences aim as if from a subject toward an object, where the "object" in this case might be a generalized bodily region (as when the infant experiences discomfort). If so, however, it is difficult to see how an infant could have a conscious experience in the absence of a concomitant sense of itself as having that experience.

There is evidence from developmental psychology that appears to support the existence of a minimal conception of self in very young infants. One experiment, conducted by Field (1976) and cited by Bermúdez (1998), involves the reaching behaviour of fifteen-week-old infants. Apparently, there is a significant behavioural difference between an infant's attitude toward an object that is within its

45. In dishonour of Hegel – who, in his dissertation *De Obitis Planetarum*, criticized Newton, gave an *a priori* defence of Kepler, and offered an *a priori* argument to the effect that there is no planet between Mars and Jupiter – this is John Baker's name for "a family of philosophical sins, membership in the family being shown by the sinner's willingness to treat what may well be empirical theses as if they were *a priori* and to do so without any proof that they are."

reach and an object that is outside its reach. When an object is placed outside its reach, there is a marked reduction in the frequency with which the infant will extend its arm toward it. Thus, the infant can discriminate between objects within its reach and objects outside its reach (Bermúdez 1998, p. 127). Bermúdez also mentions an earlier study by Bower (1972) that suggests that the ability to make such a discrimination is in place as early as two weeks. But notice that an infant can make such a discrimination only if it is aware of the position of the object in relation to *itself*, i.e., only if it possesses what is called *self-specifying information*. On the basis of this evidence, it is reasonable to conclude that very young infants have a minimal conception of self.

Notice that I have been careful to speak of a *conception* of self or a *sense* of self. All of us, I assume, have this sense of self, a sense that we are abiding (flesh and blood) creatures that undergo various experiences. But this is not to suggest that we have *selves*. Conscious creatures have various physical and psychological properties, and one of these latter is a sense of self. But we do not have selves, where having a self means that one is in possession of some kind of *inner psychological object*. This kind of alleged entity is, I submit, a philosophical fiction.

6.3 Objections and replies

In this section I canvass a handful of objections that might be raised against the modified Brentanian account of consciousness I have been defending. The objections I confront are not directed at MB as such (which is an account of the *content* of implicit self-awareness), but rather at the idea that perceiving *x* (thinking about *x*, etc.) involves a single token state that is non-objectually self-given. Nevertheless, I will use "MB" as a general term for the position under attack. The context, I trust, will make clear what is being intended.

1. A common suspicion is that since it treats consciousness as an intrinsic property, MB cannot give an informative explanation of consciousness. As I remarked (ch. 5, §5.2), Rosenthal prefers to treat consciousness as non-intrinsic for just this reason. According to Rosenthal, consciousness can be informatively explained only on the assumption that it is an extrinsic property of certain mental states. To recapitulate Rosenthal's argument: if consciousness can be informatively explained, then it must have some sort of articulable structure. But if consciousness is an intrinsic property, then it does not have an articulable structure, i.e., it is "simple and unanalyzable" (1990, p. 736). Therefore, if consciousness is an intrinsic property, it cannot be informatively explained.

Though formally valid, however, Rosenthal's reasoning is unsound: he errs in supposing that consciousness does not have an articulable structure if it is an intrinsic property of certain mental states.[46] According to MB, consciousness is an intrinsic property, but it also has an articulable and therefore informative structure.[47] To wit, a conscious state, unlike an unconscious state, is immediately and non-objectually given in experience – this is what *makes* the mental state conscious[48] – which means that, far from being simple and unstructured, a conscious state is a complex state that has a definite structure. MB, then, provides just the sort of informative explanation of what constitutes a conscious state that Rosenthal is seeking, but it does so without incurring the pitfalls of higher-orderism.

2. Another (related) objection, sometimes voiced by those seeking a reductive account of consciousness, is that as an "account" of state consciousness MB is undesirably elusive. MB says that a mental state is conscious (at *t*) if and only if the subject of the mental state is implicitly aware of it (at *t*) such that she is immediately and non-objectually acquainted with it as hers. But it might be objected that it is entirely unclear how a subject can become acquainted with a mental state in the appropriate way.

This is an imaginary difficulty. It confuses an account of the experiential structure of conscious states, which MB is, with an account of the underlying mechanisms that subserve that structure, which MB is not. MB is not, and does not purport to be, a scientific explanation of state consciousness, i.e., it does not offer an account of the underlying physical (presumably neural) mechanisms or laws governing state consciousness. Rather, as I have mentioned, it offers an answer to the preliminary question of what a mental state's being conscious consists in.

---

46. Also see Caston 2002, pp. 779–80; Gennaro 1996, p. 23; and Schröder 2001, pp. 33–4.

47. But then this is something Rosenthal himself admits. He says that Brentano's model of consciousness is a "rare attempt to assign informative structure to state consciousness, conceived of as an intrinsic property" (1990, p. 749n24). But then he rejects the Brentanian model on other grounds, namely, that Brentano "gives no reason to sustain his insistence that this awareness of conscious mental states is intrinsic to those states" (ibid.).

48. Gennaro (2006, pp. 226–8) has argued that Brentano's account of consciousness is deficient because it does not provide an *explanation* of what constitutes a conscious state. According to Gennaro, Brentano's account *describes* the structure of conscious states, but it is silent on the crucial question of what *makes* a mental state conscious. This difficulty does not attend MB, even if it attends Brentano's original account (though I am not convinced of this). Under MB, a mental state becomes conscious when it acquires a particular sort of experiential self-givenness, i.e., when the subject of the mental state becomes non-objectually acquainted with it. This is not just a description of the structure of the mental state, but also an explanation of what *makes* it conscious.

As such, the onus is not a defender of MB to articulate the sub-personal mechanisms that underlie consciousness. All a proponent of the view is required to do is give plausible grounds for supposing that state consciousness is best understood as involving a non-objectifying self-givenness; he need no articulate the precise mechanisms that subserve that structure. Indeed, to articulate the mechanisms that subserve that structure would be to articulate a solution to what Seager calls the "generation problem," the problem of explaining how complex configurations of insentient physical particles, processes, etc. can give rise to (subvene, cause, realize, whatever) consciousness.[49] But nobody has the faintest idea of what an adequate solution to this problem might look like.[50] To be sure, a few intrepid theorists have offered connectionist, neurocomputational, or quantum mechanical accounts of how bundles of neurons might achieve certain experiential effects (e.g., Churchland 1995; Crick and Koch 1990; Penrose 1989), but generally even these theorists concede that our knowledge of the physical mechanisms underpinning consciousness is highly speculative and conjectural at best.

3. The third objection is that MB is committed to the notorious Cartesian thesis, famously derided by Ryle (1949, pp. 158–9), that we have infallible access to our own conscious states.[51] MB is, I concede, committed to an infallibility thesis of a roughly Cartesian flavour, but there is nothing amiss about this infallibility thesis, and it has no dubious epistemic consequences. According to MB, subjects are immune to error with respect to their own conscious states in the following sense: whenever a subject is in a conscious state, the subject knows he or she is in that conscious state. This is an immediate consequence of MB: token conscious states are immediately given in experience, which means that there is no distance between a token conscious state and one's awareness of being in that state. Thus, if a conscious state occurs it cannot fail to be manifested in experience. But if a

---

49. See Seager (1999, pp. 17–32) for a succinct exposition of the generation problem. The problem goes by various names: Chalmers (1996) calls it the "hard problem" and Levine (2001) calls it the problem of the "explanatory gap."

50. Bennett and Hacker (2003, pp. 302–7) have argued that the generation problem is a philosophical pseudo problem engendered by conceptual confusion. Perhaps they are right. However, even if they are right, their arguments (assertions?) do nothing, it seems to me, to disentangle whatever conceptual entanglements are causing the "problem." See Seager (1999, ch. 1) for a forceful argument against this sort of attempt to dissolve the generation problem.

51. Philosophers these days tend to pillory or summarily dismiss any view that exhibits even a hint Cartesianism. According to Seager (1999, ch. 8), this tendency has resulted in a philosophical catastrophe of sorts, namely, the "anti-Cartesian catastrophe."

subject is in a conscious state that is manifested in experience, then the subject cannot fail to know that he or she is in that conscious state.

Let me examine this reply more closely, drawing on Hossack's recent analysis of consciousness and its relation to self-knowledge.[52] Let C be an arbitrary conscious state, let s be some subject, let 'Cs' designate that s is in C, and let 'KsCs' designate that s knows she is in c. All can agree that, necessarily, s knows she is in C only if s is in C. That is:

$$\Box(KsCs \rightarrow Cs)$$

Since knowledge is factive, this much is uncontentious.[53] However, as I have said, according to MB it is also the case that, necessarily, s knows that she is in C *if* s is in C. That is:

$$\Box(Cs \rightarrow KsCs)$$

It follows that:

$$\Box(KsCs \leftrightarrow Cs)$$

But then this seems eminently tenable, since one could hardly be looking at something and not know it, be hearing something and not know it, etc.

It does not follow, however, that one cannot be mistaken about the *object* or, more generally, the *content* of one's experience. It is obviously true that sometimes we err in our perceptual judgements, thinking we saw (heard, felt, etc.) something when we saw (heard, felt, etc.) something else. And sometimes we err in our emotional judgements, thinking we are feeling a particular emotion, say anxiety, when we are feeling something quite different, say anger. But all that follows from this is that sometimes we lack the information necessary to make a sound judgement about the object or content of our experience, or that sometimes our perceptual and emotional capacities at the experiential level outstrip our conceptual/cognitive resources. It does not follow, and it is false, that one can be mistaken about

52. See Hossack 2002, pp. 170–1. I have altered Hossack's logical notation.

53. There is a superficially similar (though importantly different) thesis that *is* contentious, namely, that first-person *judgements* or *beliefs* about one's own conscious states are authoritative or infallible. According to this thesis, s *believes* she is in C only if s is in C. But this is disputable. If s is having an hallucinatory experience of *x*, for example, S may believe she is veridically perceiving *x*, but surely S's *belief* that she is veridically perceiving *x* does not suffice for S to be veridically perceiving *x*. The following reminder from Hossack is to the point: "since there is no limit to human folly, the fact that someone believes something is no reason to suppose that what he believes is true" (2002, p. 176).

being in a conscious state of some sort.[54] To illustrate this point, Hossack invokes the *de re* formulation of knowledge attribution, according to which "a knowledge attribution. . . is correct iff *S* knows the proposition expressed by the content clause under some mode of presentation or other" (Hossack 2002, p. 175n16). Thus, "it is not being claimed," says Hossack, "that if one hallucinates a dagger, one knows that one hallucinates a dagger. It is claimed only that one is aware of the hallucination, i.e., one knows *de re* of one's hallucinating a dagger" (2002, p. 175).[55] Examples are easily multiplied. I can be looking at a stick water and mistakenly think it is bent, but I can hardly be looking at a stick water and not know, *in my experience* (i.e., *de re*), that I am looking at something, whatever it is. Likewise, I can be feeling a particular bodily sensation, say a pain, and not know what pain is. But an air of unintelligibility attends the supposition that one can be consciously experiencing pain and not know, in one's experience, that one is feeling pain.[56]

It turns out, then, that MB can parry a number of standard objections. This is an important result, for, in Chapter Eight, I make the claim that phenomenal character is constituted by implicit self-awareness, a claim the truth of which depends, ultimately, on the existence of a logically coherent and phenomenologically sensitive model of consciousness that takes all conscious states to necessarily involve an implicit self-awareness.

Before advancing my positive account of phenomenal character, however, it will be useful to assess the virtues and shortcomings of an approach to the analysis of phenomenal character that has received considerable attention recently, namely, representationalism. I shall try to do this in the following chapter. The representational theory of phenomenal character is directly opposed to the analysis of phenomenal character I seek to defend, and so a closer look at the theory will serve as a convenient lead up to an exposition of my account.

---

54. As Caston has noted: "[one] need only maintain that we are aware (veridically) that each token activity takes place and that it have a certain content" (2002, p. 798).

55. It is no objection to insist that there are no hallucinatory daggers. The exact ontological status of the content of hallucinations, delusions, after-images, etc. is one of the more recalcitrant issues in discussions of intentionality, but it is plausible to claim that, in having an hallucination, a subject hallucinates *something*. Dretske (1999) has argued that uninstantiated universals comprise the content of hallucinations. This might strike some as outlandish, but there is nothing obviously wrong with the idea.

56. Cf. Thomasson 2000, p. 196; Caston 2002, pp. 798–9; and Shoemaker 1996b, pp. 226–9.

CHAPTER 7

Representationalism

7.1 Preamble

In this chapter I provide an assessment of an influential extant theory of phenomenal character, viz., representationalism. There are at least three important reasons for doing this: (1) several leading philosophers of mind, most notably Tye (1995, 2000) and Dretske (1995, 2003),[1] have recently defended versions of the view; (2) the arguments in support of representationalism require independent rebuttal – I have not saved myself this labour by virtue of anything I have said thus far; and (3) the debate over whether some kind of representational account of phenomenal character is plausible is far from settled; in fact, Block has referred to the division between proponents and opponents of representationalism as "[t]he greatest chasm in the philosophy of mind – maybe even all of philosophy" (1996, p. 19). Like it or not, then, it seems that representationalism cannot be ignored: one must work within the theory or explain why not.

7.2 The representational theory of phenomenal character

In this section I raise some objections to the representational theory of phenomenal character through an examination of Dretske and especially Tye's accounts. Though similar in many ways, Tye's theory is more comprehensive than Dretske's, so I will take Tye's views to be representative of the position. Nevertheless, I will bring in Dretske where he might be required to improve on Tye or where he addresses points not considered by Tye. It is impossible within the scope of this book to examine all the variants of the theory, but a critique of Tye and Dretske's accounts, given the authoritative status usually accorded to them, should provide

---

1. See also Lycan 1996 and Rey 1998. Byrne (2001), Crane (2003), Harman (1990), Jackson (2004), and Shoemaker (1994, 1996c), also defend representationalism, though the kind of representationalism they defend – what might be called "weak" or "supervenience" representationalism (see below in this section) – is not my interest in this chapter.

good grounds for supposing all versions are unsound. I shall argue that the representational theory of phenomenal character, at least on Tye and Dretske's analysis, rests on at least two unsupported and erroneous – i.e., phenomenologically implausible – assumptions.

Representationalism is the thesis that, as Dretske puts it, "all mental facts are representational facts" (1995, p. xiii). The notion of "representation" at work here is not always clear, but representationalists typically equate representation with intentionality.[2] The idea is that insofar as our conscious experiences are *of* or *about* objects, properties, events, etc., our experiences *represent* those objects, properties, events, etc. When representationalists say conscious experiences are representational, however, they are saying more than that conscious experiences are intentional, i.e., of or about something. Representationalism nowadays is tied to a naturalizing project, whereby conscious experiences are to be *explicated* in terms of mental representations, and these mental representations are complex *information-providing* entities housed in the brain. Thus, when representationalists use the term "representation," they use it as roughly synonymous with "indication" or "stand for." On this view, conscious experiences are physical (presumably neurophysiological) states that, just like tree rings, columns of smoke, speedometers, tachometers, etc., indicate or convey information about some putative fact or property. So just as a tachometer conveys information about an engine's revolutions per minute, so conscious experiences convey information about various perceptible properties like size, shape, colour, texture, and so forth.[3] The guiding assumption is that explaining

2. Cf. the following comments by, respectively, Tye and Rey: "Brentano did succeed in drawing to our attention an extremely interesting fact about a wide range of our mental states (whether or not they are all ultimately like this), namely, their capacity to represent or be about things that do not exist" (Tye 1995, p. 95); "[a]n increasing number of philosophers…have proposed 'Representationalist' theories of the qualitative character of especially sensory experience, according to which the apparent objects and properties of that experience are merely 'representational' or 'intentional' (I will use these terms interchangeably)" (Rey 1998, p. 435).

3. This is, I think, an accurate statement of the representational theory, at least as expounded by Dretske and Tye. However, the claim that our experiences represent – in the sense of convey information about – properties in our environment is far from unpuzzling. Seager argues that the representational theory is "phenomenologically verified" (1999, p. 134). "I irresistibly regard," he claims "perceptual experience as informing me about the local environment" (ibid.). This way of characterizing the theory, however, is bound to engender confusion, for it is far from clear whether our perceptual experiences "inform" us about anything. If I look about the lecture hall, I might see a lectern, but my seeing the lectern does not inform me about the presence of the lectern. Bennett and Hacker put the point this way: "It is misleading to say that the senses are concerned with 'telling us what is there'.… For the

consciousness in terms of representational states naturalistically underpins consciousness. This is because the world is filled with natural systems of representation, and if conscious states are equally representational, then, *ex hypothesi*, conscious states are equally natural.

It should be pointed out that representationalists are quick to distinguish between representational *vehicles* and representational *contents*. What are in the heads of subjects are representational vehicles, not representational contents. As Dretske says:

> We can have stories about blue dogs (blue-dog stories) and blue-dog experiences. Stories about blue dogs – the blue dog vehicles – are neither blue nor doglike. Just look in the book. It is all black and white. Just so, what we find by looking in the brain of a person experiencing blue dogs is neither blue nor doglike. We do not find the content of experience, the properties that make the experience the kind of experience it is. What we find, instead, are electrical and chemical activity in gray, cheesy brain matter. We find the experience vehicles.
> (1995, p. 36; cf. Lycan 1990, p. 111)

The point is straightforward, the analogy persuasive. In our brains we have representational vehicles that represent or indicate things like blueness, roundness, etc., but there is no blueness or roundness in our brains. It is just a matter of our brains "telling us" or indicating to us that something is blue or round without these properties being in our brains.

It should also be emphasized that although representationalists use the term "representation" as roughly synonymous with "indication" or "stand for," there are important differences between these terms depending on which general conception of representation underlies one's representationalist view. Dretske, for example, opts for a teleo-evolutionary semantics, arguing that the relevant criterion for understanding representation is *proper function*, whereas Tye opts for a less constrained conception of representation, arguing that the relevant criterion for understanding representation is "correlation" or "causal covariation," *under optimal conditions* (which he calls *tracking*). Thus, for Dretske, "there is no representation without functions" (1995, p. 4), which means that tree rings, though they convey information about the age of

---

senses are not information-transmitters. Our eyes do not see things and tell the brain what they see, and the brain does not tell us what the eyes see. We use our sense-organs to find out, observe, apprehend what is in our environment. Our sense-faculties are capacities for apprehending, by the use of our sense-organs, how things are in our environment. The senses are not information-transmitters, although we acquire information by their use" (2003, p. 146). There is reason to suppose, then, that the representational theory is not even coherently stateable. However, for present purposes, and in the spirit of charity, I will grant that it is.

a tree, cannot be said to represent the age of a tree, since that is not their function. For Tye, by contrast, tree rings do represent the age of a tree, since, under the causal covariation conception of representation, the power to represent is independent of the representing object's functional design. The details of these different conceptions of representation need not delay us; the objections I level against representationalism will not depend on resolving the debate over which is preferable.

Now if representationalism is true, then all mental facts are representational facts. It follows that the fact that our experiences have phenomenal character is a representational fact, since the fact that our experiences have phenomenal character is a mental fact. But what kind of representational fact could the fact that our experiences have phenomenal character be? According to Tye and Dretske, phenomenal character is a species of representational *content*, and, on Tye's theory, a *specific kind* of representational content, namely: Poised, Abstract, Nonconceptual, Intentional Content (hence the acronym PANIC Tye uses to describe his theory; see Tye 1995, ch. 5 and 2000, ch. 3).[4] When one perceives an object one's perceptual experience represents that object as possessing certain properties (e.g., blueness, roundness, etc.). These properties comprise the representational content of the experience, and the phenomenal character of the experience is *constituted by* this representational content (hereafter simply "content"). As Dretske puts it: "experienced qualities, the way things phenomenally seem to be (when, for example, one sees or hallucinates an orange pumpkin), are – *all* of them – properties the experience represents things as having" (2003, p. 67). This means that the properties a perceived object is represented as having, whether one's perception of the object is veridical or not, exhaust what it is like for a subject to perceive that object. Thus, if I have a visual experience of blue (say I am looking at a clear blue sky), the properties this inner state represents, viz., blueness, exhaust the phenomenal character of my visual experience of blue. About this, Dretske says:

> On almost any theory of experience – certainly on a representational theory – there is no more to the quality of one's experiences in experiencing blue than there is to the color blue since the color blue *is* the color one experiences. It is the quality that distinguishes one's blue experiences from other color experiences.
>
> (1995, p. 85)

Thus, the representational theory is able to account for differences in phenomenal character in representational terms because it identifies phenomenal character with content. The phenomenal character of an experience changes when its content changes.

---

4. I will not dwell on the details of Tye's PANIC theory. My critique of representationalism does not depend on resolving any controversies about how to characterize the representational content with which phenomenal character is supposedly identical.

There are weaker forms of representationalism; for example, one might argue that phenomenal character *supervenes* on content, but is not reducible to it. On this weaker version of the theory, content determines, or provides the subvenient base for, phenomenal character, such that, necessarily, "there can be no difference in phenomenal character without a difference in content" (Byrne 2001, p. 204). Thus, if two experiences differ in phenomenal character, then they differ in content. This weaker version is silent, however, on the nature of phenomenal character, and so representationalism's prime exponents adhere to a stronger version, one that, as Tye puts it, "aims to tell us what phenomenal character *is*" (2000, p. 45). On this stronger version, content fully exhausts the nature of phenomenal character. Since it purports to explain the nature of phenomenal character, and is thereby in competition with the theory of phenomenal character I propound in Chapter Eight, this version of the theory will be the focus of my critique.

To show precisely wherein this version of the representational theory of phenomenal character falters, I will provide an extended critique of two of the theory's central motivating assumptions. I will call these assumptions the Transparency Assumption and the Property Assumption, respectively.

7.2.1 The Transparency Assumption

The recent proliferation of representational theories of consciousness is attributable, at least in part, to an assumption that even critics of representationalism often unhesitatingly accept (see, e.g., Rosenthal 1990, p. 744; Carruthers 2000, p. 183; and Rowlands 2001, p. 161), namely, that conscious experiences are *diaphanous* or *transparent*. The Transparency Assumption is, in effect, an allegedly uncontroversial phenomenological observation (see Shoemaker 1994, pp. 30–1). When we attempt, it is claimed, to introspectively focus on the qualitative features of a conscious experience, we will always find ourselves focusing our attention on the features of the object of the experience, not on some intrinsic feature of the experience itself. If we try to focus more intently on some intrinsic feature of the experience itself, on some feature of the experience that distinguishes it from other experiences, we just end up attending more closely to the features of the object we are experiencing. This supposedly shows that conscious experiences are transparent, i.e., that there is no more to the quality of an experience than the qualities of the object one is experiencing.[5] Consider the following oft-quoted remark by Harman:

---

5. Moore is usually cited as having diagnosed the problem of transparency. In a famous essay he writes: "And, in general, that which makes the sensation of blue a mental facts seems to escape us: it seems, if I may use a metaphor, to be transparent-we look through it and see nothing but blue" (1922, p. 20).

> When you see a tree, you do not experience any features as intrinsic features of your experience. Look at a tree and try to turn your attention to intrinsic features of your visual experience. I predict you will find that the only features there to turn your attention to will be features of the presented tree, including relational features of the tree "from here."
>
> (1990, p. 39)

Tye agrees:

> Try to focus your attention on some intrinsic feature of the experience [of a blue square] that distinguishes it from other experiences, something other than what it is an experience *of*. The task seems impossible: one's awareness seems to slip through the experience to blueness and squareness, as instantiated together in an external object. . . . [I]ntrospection does not seem to reveal any further distinctive features of the experience over and above what one experiences.
>
> (1995, p. 30)

Proponents of representationalism argue that the transparency of experience leads naturally to the view that phenomenal character and content are one and the same, i.e., that there is no more to the quality of one's experience of *x* than the qualities one's experience of *x* represents *x* as having.

Before proceeding any further, I want to make a distinction the lack of which might muddy the discussion to follow, the distinction, namely, between *strong* transparency (hereafter ST) and *weak* transparency (hereafter WT). ST and WT may be formulated as follows:

(ST) It is impossible to discern, via introspection, any intrinsic properties of an experience of *x* that are not experienced as properties of *x*.[6]

(WT) It is difficult, but not impossible, to discern, via introspection, intrinsic properties of an experience of *x* that are not experienced as properties of *x*.

Distinguishing these two versions of the Transparency Assumption is important because ST, not WT, motivates the representational theory of phenomenal character that Tye and Dretske endorse.[7] As previously remarked, this version of the theory says that phenomenal character is constituted by content. But if our experiences exhibit qualitative features that are not experienced as identical to the features of the objects we experience – i.e., if ST is false – then phenomenal character

6. When Tye and other proponents of transparency say that it is impossible to discern intrinsic features of one's experience, I take them to mean that it is logically impossible.

7. A point also made by Kind (2003, pp. 238–9).

cannot be constituted by content, since content is a direct causal product, as it were, of the features of the objects we experience, in the sense that the presence of these features "activates" the internal sensory processes that generate content.[8]

Drawing a distinction between ST and WT is also important because exponents of ST sometimes present the case for ST in contradistinction to the case for sense-data (e.g., Harman 1990, p. 39; Tye 2000, pp. 45–6). The now largely defunct and unlamented (though see Casullo 1987 and H. Robinson 1994 for a relatively recent defences) sense-datum theory of perception is, roughly, a *mediational* theory of perception, according to which mental entities mediate our perception of extra-mental objects. But enthusiasts for ST argue that introspection does not reveal any of these alleged mental entities; in fact, introspection, they tell us, does not reveal any distinctive features of an experience over and above the features of the object one experiences. From this (among other things) they conclude that the sense-datum theory is probably false and that ST is true. Now the principal claim of this section is that ST is false, but this should not be construed as a tacit espousal of sense-data. My own view, in fact, is that WT is true, but WT, like ST, is entirely consistent with, and in fact strongly supports, the deep-seated intuition that, in ordinary perception at least, we are *directly* acquainted with the objects we perceive, i.e., that there is no mental intermediary between us and the objects of which we are perceptually conscious.

Kind (2003) has also drawn a distinction between strong and weak transparency, which corresponds roughly to the distinction I have drawn. However, Kind's characterization of strong transparency is somewhat misleading. She characterizes strong transparency as the thesis that "it is impossible to attend directly to our experience, i.e., we cannot attend to our experience except by attending to the objects represented by that experience" (2003, p. 230). But this seems false. It is relatively easy to shift one's attention from the object of one's visual experience, say, to one's visual experience itself and attend exclusively to features of one's visual experience. For example, if I shift my attention from the words on this page (which I am visually experiencing) to my visual experience itself, I can effortlessly discern that it is quite unlike my occurrent auditory experience of the music of Rafael Toral in the background. Similarly, if I shift my attention from the music of Rafael Toral in the background (which I am auditorily experiencing) to my auditory experience itself, I can equally effortlessly discern that it is quite unlike my occurrent olfactory

8. Cf. the following claim from Lycan: "Now, how does phenomenal color depend upon real color? My answer is, by representation: for a visual state to have or contain or feature a green quale in my strict sense is for it to represent greenness – real physical greenness – at some location in the visual field" (1996, p. 74).

experience of the tea in the mug beside my computer. And so on. This property of my visual experience, however, i.e., the property of being unlike my auditory experience, is an *extrinsic* or *relational* property of my visual experience, and so does not threaten ST. ST does not say that it is impossible to discern, via introspection, any extrinsic features of one's experience of *x* that are not experienced as features of *x*. Indeed, as the above examples illustrate, introspection reveals any number of these. To reiterate the point with emphasis, then, ST is the thesis that it is impossible to discern, via introspection, any *intrinsic* features of one's experience of *x* that are not experienced as features of *x*.

So stated, however, ST might still seem vulnerable to a number of straightforward counterexamples. Consider blurry vision, i.e., seeing something blurrily (as opposed to seeing a blurry image).[9] Blurriness is, incontestably it seems, not a feature of the object of one's visual experience, but rather a distinctive and intrinsic feature of the visual experience itself. It is a feature not of *what* is experienced, but of *the way in which* something is experienced. Or consider a person whose auditory experiences are muffled because of an inner ear infection. One could say, not implausibly, that this muffled quality is an intrinsic property of the auditory experience itself, not a property of the object being experienced. These examples alone, if accepted, indicate that perceptual experiences are not transparent in the way representationalists suggest.

Tye claims, however, that blurry vision presents no obstacle to ST. On his view, when one unfocuses one's eyes or removes one's eyeglasses, one simply "loses information." "In seeing blurrily," he says,

> one undergoes sensory representations that fail to specify just where the boundaries and contours are. Some information that was present with eyes focused is now missing. In particular, the grouped array contains less definite information about surface depth, orientation, contours, and so forth.
>
> (2000, p. 80)

It is unclear, however, whether this response answers the charge that blurry vision constitutes a counterexample to ST. The informational impoverishment of which Tye speaks is, surely, a property of one's visual experience, not a property of the object being experienced. I take it as beyond serious dispute that the actual

---

9. Tye, in his explication of the appeal to transparency (2000, pp. 45–51), concentrates mainly on visual experience, arguing that his points generalize to other perceptual modalities and, indeed, to the phenomenal character of felt moods and emotions. In my critique of the Transparency and Property Assumptions, I will, for the most part, follow Tye and focus my examples on visual experience.

extra-mental objects in the grouped array do not contain less definite information about surface depth, orientation, contours, and so forth. But if blurriness is a property of one's visual experience, then introspection does, it seems, reveal a possible feature of one's experience that is not a feature of the object being experienced.

Matters are more complicated, however. Tye would agree that blurriness is, in some sense, a property of one's visual experience, but he would deny that this entails the falsity of ST. When one sees blurrily, the qualities of which one is directly aware are all *experienced* as being qualities of the objects of one's experience, even though they are not (see Tye 1995, pp. 157–8). This means that blurriness is not revealed through introspection, but rather through inference. One infers, perhaps on the basis of formerly seeing clearly, that one is presently seeing blurrily. So blurriness does not vitiate the appeal to transparency, for ST says that it is impossible to *introspect* an intrinsic feature of one's occurrent visual experience; it does not say that it is impossible to *infer* an intrinsic feature of one's occurrent visual experience.

So Tye has a plausible response to the objection from blurry vision, a response that would work equally well against the objection from muffled audition. In fact, Tye has confronted a number of proposed counterexamples to ST (2000, ch. 4), skillfully demonstrating that the thesis can emerge unscathed in each case.

I want to propose a pair of counterexamples to ST, however, that are not as easily answered. Before doing this, let me state more clearly the attributes a successful counterexample to ST must possess. ST says that it is impossible to discern, via introspection, any intrinsic features of an experience of *x* that are not experienced as features of *x*. This means that a successful counterexample to ST must be (1) intrinsic and (2) discernible through introspection. The counterexamples I want to propose possess both these properties.

The first counterexample draws on a result from a previous chapter. Recall that, in Chapter Six, it was argued that conscious states are always accompanied by an implicit self-awareness. In seeing *x* (hearing *x*, etc.) I am implicitly aware of seeing *x*. Furthermore, I am not merely implicitly aware of *seeing x*; rather, I am implicitly aware of *myself seeing x*. All conscious experiences have this *de se* component; that is, they are all experienced as occurring, to use some quasi-Jamesian terminology, "in the experiencer's stream." This sense of self-presence is not a trivial feature of consciousness, accurately imputed to some, but not all, conscious states; rather, it is an artifact of the reflexive nature of state consciousness generally.

Let me demonstrate how the characteristic implicit self-awareness involved in conscious states satisfies the criteria for a successful counterexample to ST. Notice, first, that this feature of experience is intrinsic. In seeing *x* (hearing *x*, etc.), I am implicitly aware of myself seeing *x*, and this implicit self-awareness is built into my very act of seeing. Observe, as well, that this feature of experience is readily

introspectible. Here I would ask the reader of these words to shift his attention from the object of his act of looking to his act of looking itself, and challenge him to deny that he did not experience, however subtly, his act of looking as occurring in his stream. One might be tempted to object that, in being aware of oneself seeing (hearing, etc.), one is not *consciously aware* of oneself seeing, but rather *unconsciously aware* of oneself seeing; however, as I have already argued, this is implausible: the mental states in virtue of which we are aware of our conscious experiences are themselves conscious. One might also be inclined to object that this feature of experience, given its elusiveness, is non-introspectible. If one tries to reflect on one's implicit awareness of one's conscious state, the implicit awareness "shifts" to one's act of reflection, rendering implicit self-awareness permanently unavailable to introspection. This objection, however, relies on an exceedingly narrow view of introspection, according to which it is possible to reflect on a conscious state *c* only if one is *in c*. It is true that one cannot attend to one's implicit awareness of one's *occurrent* consciousness of *x*, but surely one can attend to, and reflect on, one's implicit awareness of one's *non-occurrent* consciousness of *x*. After all, philosophers since Aristotle have been developing theories of consciousness according to which consciousness is explicable in terms of implicit self-awareness, and the construction of such theories would seem to require that the theorist recognize and reflect on her implicit self-awareness. To be sure, such reflection is indirect, since it is mediated by memory, but there is no reason to suppose that the results of indirect introspection should be discarded. If I experience a minor ache in my neck, I may not reflect on the pain while I am experiencing it. But if I reflect on it later, after it has passed, the results of my reflection are still theoretically exploitable.

The *temporal* properties involved in conscious states constitute another counterexample to ST. Many of our conscious states have temporal properties. A pain, for example, cannot be said to be an ache unless it has duration. We refer to momentary pains as cramps, twinges, spasms, etc. Similarly, one cannot be said to be in a state of attentive thinking or concentration unless one cogitates for a period of time. It makes scant sense to say of a subject that she was in a state of intense concentration for two seconds.

These temporal properties satisfy the criteria for a successful counterexample to ST. The duration of an ache is an introspectible property of the ache; one can reflect on the chronic ache in one's lower back, for example, pondering its unremitting nature. In addition, the duration of an ache is an intrinsic feature of the ache. Similarly, even though it is typically impossible to reflect on one's occurrent state of attentively thinking about *x* (since, given the nature of focused cogitation, taking notice of and reflecting on one's occurrent state of attentively thinking about *x* typically renders one incapable of attentively think about *x*), it is possible

to reflect on one's *non-occurrent* state of attentively thinking about *x*; for instance, it is possible to reflect on the duration of one's non-occurrent state of methodically thinking about Fermat's enigma. What is more, the duration of one's state of attentive thinking is an intrinsic feature of that state of thinking.

I conclude that implicit self-awareness and temporal properties constitute counterexamples to ST.

7.2.2 The Property Assumption

The foregoing counterexamples to ST provide ample reason to be sceptical of the representationalist's attempt to reduce phenomenal character to representational content, but the representational theory of phenomenal character is criticizable on other grounds. One such ground is that it rests on what I will call the Property Assumption. This assumption says that the immediate objects of consciousness – or, more exactly, the immediate *data* of consciousness – are not objects per se, but rather *properties*, such as redness, blueness, roundness, shininess, etc. Unlike the Transparency Assumption, few representationalists have explicitly stated their avowal of the Property Assumption. As will be seen below, Tye endorses the assumption unambiguously, but it is not clear that other representationalists do, so it is perhaps somewhat infelicitous to call the Property Assumption a "motivating assumption." Nevertheless, a strong case can be made that representationalism cannot rid itself of the assumption. That is to say, even though some representationalists may not be "motivated" by the Property Assumption, they are nevertheless theoretically committed to it, in the sense that they cannot consistently cleave to representationalism and deny it.

Here is why. Denying the Property Assumption while retaining representationalism would involve demonstrating that the items represented in experience can be objects, and not merely properties of objects; in short, it would involve demonstrating that representational content can be *object-content*. But this cannot be done. To see this it is important to bear in mind what representationalists mean when they claim that our sensory systems represent things. As I mentioned above, the notion of representation, at least as it is employed by representationalists seeking to reduce phenomenal character to representational content, is, minimally, a causal notion.[10] The idea, baldly stated, is that there is an evolutionarily conferred tracking relation between states of a creature's perceptual system and various features in the creature's environment that causally generates the content with which phenomenal character is purportedly identical. But what sorts of features in the environment are involved in this tracking relation? It is quite clear that the relevant features cannot be whole

10. See Dretske 1994, p. 133.

objects, such as pens, watches, chairs, pianos, pop bottles, etc., since to affirm that, say, a pop bottle is involved in generating content would be to deny, implausibly, that percipient creatures can visually experience objects (like pop bottles) that played no part in the evolutionary development of the creature's perceptual system. For presumably there are no states of a creature's perceptual system that have evolved to causally co-vary specifically with pop bottles. So, assuming that we have acquired our complex perceptual systems evolutionarily and that these systems enable us to consciously perceive things like pop bottles, it cannot be the pop bottle that is involved in generating content, but rather the properties of the pop bottle, i.e., its colour, its cylindricality, etc. But if the properties of the pop bottle are involved in generating content, then the items represented in one's experience of the pop bottle are, in the first place, the properties of the pop bottle, not the pop bottle itself.

It would seem, then, that representationalists are committed to the Property Assumption, whether or not they are in fact motivated by it. In any case, Tye endorses the assumption unequivocally, and, as suggested by the following remark in which he portrays the assumption as a phenomenological datum, it is reasonable to suppose that he is motivated by it as well:[11]

> In general, we see things by seeing their facing surfaces, and we see facing surfaces by seeing their color. Intuitively, there is no more basic level of seeing than the seeing of colors – just as there is no more basic level of hearing than the hearing of loudness, pitch, and so forth.
>
> (2000, p. 101)

---

11. The Property Assumption also motivates Dretske and Lycan. On Dretske's view, experiences represent facts, where a fact is an instantiation of a property by an object. But if our experiences represent properties instantiated by objects, then we experience (at least in the first place) properties, not objects. The following remarks reveal Dretske's adherence to the Property Assumption (my italics throughout): "These differences [between sensory and conceptual representations] are important in thinking about ways of representing properties and, thus, according to representational theory, *the properties we are conscious of*" (Dretske 1995, p. 11); "What, then, in representational terms, is the basis of this ordinary…distinction between an *experience of color, shape, and texture*, and a belief or judgement about color, shape, and texture?" (Dretske 1995, p. 12); "There are two…compelling reasons for thinking of experiences in representational terms. There is, first, the fact that the *properties that individuate experiences*, the ones that distinguish one of type of experience from another…are not (at least they need not be) qualities of anything in the experiencer…" (Dretske 2003, p. 70). And on Lycan's view, our experiences represent (at least in the first place) properties. Lycan often speaks of having sensory experiences of red, blue, etc. (e.g., 1996, pp. 69, 76), and he adopts the "placeholder" view that "an apparent visual object is phenomenally green, and its containing sensation is 'of green' or represents greenness, iff this sort of sensation is 'normally caused by' green physical objects" (1996, p. 75).

But straight away something is amiss. When we look around and see things we do not usually attend to the colours of the objects we see. Colours, though always *part* of the content of our visual percepts (since the objects we see come coloured), do not, by and large, form the intentional object of our visual experiences. We can, of course, attend to the colour of something – the colour of the paint for example – in which case the colour of the object is the object of our visual experience. But phenomenology reveals that, for the most part, visual experiences are object-involving: we see coloured objects, not the colours of objects. If I see a blue pen, then I see a blue *pen*.[12] So it seems false to say that there is no more basic level of seeing than the seeing of colours.

The dubiousness of the view that the most basic level of seeing is the seeing of colours becomes apparent in light of examples involving visual experiences of multi-coloured objects.[13] Across from my desk is a bookshelf containing a hundred or so books; and when I look at the bookshelf, that is what see: a bookshelf containing a hundred or so books. The spines of these books are red, green, blue, purple, etc. and all manner of shade of red, green, blue, purple, etc. Now Tye would have us believe that when I look at the bookshelf I have a visual sensation of some kind. He would, for example, have us believe that when I look at a particular book, say one with a blue spine, I have a bluish sensation. But what about my visual experience of the *bookshelf containing the books*? That is, what about my visual experience of the bookshelf containing the books "as a whole," i.e., as a discrete object among others? Here things get tricky. Presumably Tye would not say, *prima facie* absurdly, that I have a *bookshelf-containing-a-hundred-or-so-coloured-books* sensation. But nor can he say that I have a sensation of red, green, blue, purple, etc., since I do not attend directly to any of these colours. I see a coloured bookshelf containing a hundred or so coloured books, but I do not attend directly to any particular colour. So it cannot be the case that I have a red or green or blue or purple

12. It has been claimed that, strictly speaking, we never really see objects, only their surfaces (e.g., Broad 1925, p. 148ff.). This position I am inclined to dismiss, perhaps dogmatically, as unserious, or as more or less immediately confuted by experience. To be sure, when I see an object, I do not see all parts of it, but as Campbell has pointed out (2004, p. 400), seeing an object does not require seeing all parts of it, just as touching an object does not require touching all parts of it.

13. Non-coloured objects, such as clear water and glass, pose a problem for this view as well. Tye, in claiming that the most basic level of seeing is the seeing of colours, tacitly assumes that a thing needs to be coloured to be visible. But there is plainly room for disagreement here. What colour *is* clear water?

colour sensation. And nor can Tye say that I have a "conjunctive colour sensation," i.e., a sensation of red, green, blue, *and* purple, since I do not attend directly to an assemblage of colours in the same way I might attend directly to an assemblage of colours if I were looking at a colour-field painting. It seems that Tye would have to say that I have a visual experience of a bookshelf containing a hundred or so books, which is true. But then I do not have a colour sensation of any kind; rather, I see a bookshelf containing a hundred or so books. So where the colour sensation? Where the seeing of colours, as opposed to the seeing of a coloured thing?

Now it might be objected that my interpretation of the representational theory has been less than maximally charitable. There is an interpretation of the theory, this objection continues, according to which it is, in fact, possible to retain the theory while rejecting the Property Assumption. On this construal of the theory, bookshelves, pens, watches, chairs, pianos, pop bottles, etc. do in fact constitute the intentional *objects* of visual consciousness, but these intentional objects are composed of coloured facing surfaces that form the immediate representational *contents* of our visual experiences. The basic idea is that our visual systems generate sensory representations of coloured surfaces, and these coloured surfaces possess any number of additional properties, such as redness, squareness, roundness, rigidity, shininess, etc., that go into forming the representational content of a visual experience. Thus, if I am looking at a shiny red car, the intentional *object* of my visual experience is a shiny red car, but the *content* of my experience consists of a red surface that possesses its own properties, among them shininess, and it is these properties – i.e., the properties that form the content of my experience – with which I am directly perceptually acquainted.

This seems, in fact, to be Tye's considered position. According to Tye, "vision is modular in the sense that it generates representations of a certain class of properties of distal stimuli via perceptual processes that operate on the retinal input in a largely fixed, autonomous manner" (1995, p. 102). To buttress this hypothesis, Tye draws on the standard computational approach to vision. He writes:

> . . .representations are built up of distal features of the surfaces of external objects in mechanical fashion by computational processes. The initial or input representations for the visual module track light intensity and wavelength. . . . The output representations track features of distal stimuli under optimal or ideal perceptual conditions. Thereby. . .they represent those features, they become sensations *of* edges, ridges, colors, shapes, and so on.[14]
>
> (1995, p. 103)

14. One wonders whether Tye has not transgressed the bounds of sense here. What can it mean to say that when I look at the edge of, say, a table I have a sensation of an edge?

But there is more to the story, for one's visual field usually appears as an integrated whole, not as a random array of edges, ridges, colours, shapes, and so on. Such integration occurs, according to Tye, when specialized visual modules provide inputs to a more generalized module:

> In the case of vision, the pertinent module is the one that has as its input a unified representation of the entire visual field. . . . This is served by more specialized visual modules. Their inputs (for surface color or distance away, for example) provide inputs to the more generalized module, within which integration occurs and the final overall nonconceptual representation of surfaces and their features is constructed.
>
> (1995, p. 138)

On Tye's theory, then, visual systems generate representations of redness, squareness, roundness, rigidity, shininess, etc., and a generalized visual module integrates these representations by constructing an overall non-conceptual representation. So if I look at a shiny red car, my visual system will generate mental representations of, among other things, redness and shininess, and a more generalized visual module will integrate these representations to form an overall non-conceptual representation of a shiny red car.[15]

A number of serious difficulties attend this account of how we come to consciously perceive things in our visual field, however. Let me outline, somewhat more vividly, Tye's account of how we come to consciously perceive discrete things as discrete things, of how the sensory manifold comes to have the minimal sense of "things and the spaces between things." On the representational view, perception is not immediately the discrimination of things and the spaces between them, as reflection on our actual experience would suggest, but rather a matter of having representations that correspond to stimuli, the stimuli being *properties* of discrete things. The primary external objects of perception are discrete properties of things – colours, shapes, sizes, etc. – not discrete things or wholes. Our perception of discrete things would then seem to be a matter of the association and integration of the representations corresponding to the properties of ultimately discrete things.[16]

---

15. It should be emphasized that the generalized visual module does not integrate the actual properties of the shiny red car, but rather *representations* of these properties. This aspect of Tye's position is worth underscoring because it makes no sense to say that our perceptual systems integrate, into a perceptual experience, the actual extra-mental properties of the object of one's perceptual experience.

16. This paragraph is an almost verbatim transcription of Dwyer's (1990, pp. 15–16) characterization of classical empiricist accounts of how we come to perceive discrete things as discrete things. The tempting moral: Tye's representationalism is classical empiricism in modern guise.

To use Ryle's phrase, however, such an account "flouts the testimony of experience" (1949, p. 44), which is to say that such an account is phenomenologically idle, answering to nothing in our experience. On this, I am sure, all can agree.

The inevitable response to the charge of phenomenological idleness is to make recourse to unconscious mental activities. "Unconscious" is not meant here in any Freudian sense, of course, but simply as automatic and so unnoticed. The idea is that the representational activity is sub-personal; it takes place "behind the scenes," as it were, so we have no phenomenological access to it. Now, ordinarily, the kinds of things one does unconsciously or automatically, such as blinking, breathing, or moving one's eyes while reading, can be noticed upon reflection. But this is not the case with the unconscious mental activities proposed here. The unconscious mental activities proposed here are permanently phenomenologically inaccessible, since they occur deep within the cognitive architecture of the brain. And what phenomenology reveals is the finished product, viz., conscious experiences of bookshelves, pens, watches, chairs, pianos, etc. Our sensory systems, through some as yet unknown intellectual-computational process, generate conscious mental representations that indicate the presence of these things.

It seems to me, however, that all Tye has done here is reduced one task to another, and this is explanatorily fruitful only if we know how the second task is done. But it is altogether mysterious how a person's having access to putative representations corresponding to the properties of an object or event enables the person to *visually experience* that object or event; and saying that unconscious mental processes manipulate these representations to produce conscious visual experiences only compounds the mystery. Thus, in addition to being phenomenologically idle, the appeal to unconscious mental activities is explanatorily idle. Furthermore, the very existence of such unconscious mental activities is by no means beyond dispute, especially if we follow the lead of Wittgenstein and Merleau-Ponty, who assailed attempts to explain perception in terms of inner mental goings-on.[17]

---

17. To restate an observation made by Dwyer (1990, p. 35), the appeal to unconscious mental activities is what Wittgenstein once called "that dead-end in philosophy, where one believes that the difficulty of the task consists in our having to describe phenomena that are hard to get hold of, the present experience that slips quickly by, or something of the kind" (1953, §436). And both Wittgenstein and Merleau-Ponty were influenced by the Gestalt school of psychology (see Dwyer 1990, passim), which eschewed the notion of sensory representations that correspond to simple qualities of perceived objects. According to Gestalt psychology, the organized whole is primitive, i.e., "a figure on a background is the simplest sensible given available to us" (Merleau-Ponty 1962, p. 4).

There is another ground on which to base a rejection of the foregoing attempt to rescue the representational theory, namely, that the attempt makes the theory look suspiciously similar to an indirect realist or mediational theory of perceptual consciousness. On the proposed analysis, the immediate data of perceptual consciousness are properties of objects, and we come to perceive the object that instantiates these properties *via* or *by way of* unconscious mental processes. But, of course, this smacks of a mediational picture of perceptual consciousness of the sort from which representationalists are, generally keen to distance themselves (see, e.g., Dretske 1995, pp. 127–9; Harman 1990, p. 39; Tye 2000, pp. 45–6).

A possible objection here is that although representationalists typically maintain that the representational theory can accommodate the intuition that we are directly acquainted with the objects of perceptual consciousness, it does not follow that representationalism as such is wedded to the assumption that perceptual consciousness has an entirely unmediated structure. In other words, it might be true that, as a matter of sociological fact, the representational theory is sometimes touted as being a non-mediational theory, but one can acknowledge that the theory is, in some sense, mediational without abandoning it.

Two considerations suggest that a representationalist would be keen to resist this reply. First, if the representational theory is in fact a mediational theory, then the question arises whether it affords any serious theoretical advance over the sense-datum theory, the quintessential mediational theory of perception. Second, and relatedly, a number of representationalists have explicitly stated their opposition to mediational theories – Dretske, for example, has called the sense-datum theory "lame" (1995, p. 129) – so the onus would fall on a representationalist to demonstrate how representationalism, itself a mediational theory, can escape the multitude of objections levelled against other mediational theories.

It seems, then, that the Property Assumption is fraught with difficulties. On the one hand, the assumption is belied by ordinary perceptual consciousness: the immediate objects of perceptual experience are not, at least not always, properties of objects. On the other hand, rejecting the assumption while retaining representationalism requires invoking highly speculative cognitivist/intellectualist considerations – e.g., unconscious representational-computational processes – that, as Wittgenstein and Merleau-Ponty have shown, are eminently questionable. Moreover, invoking these cognitivist/intellectualist considerations makes the representational theory resemble an indirect realist theory of perceptual consciousness of the sort representationalists are wont to avoid.

The foregoing probings suggest that neither the Transparency Assumption nor the Property Assumption procures for the representationalist what she ultimately seeks, viz., strong support for her attempt to reduce phenomenal character to representational content. And since these assumptions are two of the pillars

LIVERPOOL JOHN MOORES UNIVERSITY
LEARNING SERVICES

on which the representational theory of phenomenal character rests, the falsity of these assumptions provides good grounds for supposing that the theory is unfounded.

There is, though, a more general moral to be extracted. Both the Transparency Assumption and the Property Assumption were shown to be phenomenologically idle, which suggests that those who would underwrite their analyses of phenomenal character with these assumptions have not done the requisite phenomenological spadework (or have not consulted the efforts of those who have, such as Brentano, Sartre, and Merleau-Ponty). The moral, then, is one Zahavi has recently pressed, namely, that "both analytical philosophy and phenomenology would profit from a more open exchange" (2002a, p. 7).

The considerations advanced in this chapter, then, show that the representational theory of phenomenal character is rejectable on phenomenological grounds. In the next chapter I develop and defend an alternative (and novel) theory of phenomenal character, according to which the phenomenal character of conscious experience is constituted by a type of implicit self-awareness. This approach to the analysis of phenomenal character is able to get round the difficulties besetting representationalism while allowing us to progress some distance toward resolving certain problems that arise in connection with phenomenal character – e.g., the problem of whether one can deny that conscious thoughts have a distinctive "feels" but nevertheless affirm that there is something it is like to have them.

CHAPTER 8

The nature of phenomenal character

> "There is less to 'the mind' than we philosophers have supposed."
>
> (Collins 1987, p. 43)

8.1 Preamble

It is generally assumed that what makes the problem of consciousness so intractable is the phenomenal character of experience; and it is also widely assumed that although self-awareness is a difficult problem in its own right, it cannot shed light on the nature of phenomenal character. The following quote from Woodruff Smith is illustrative of the tendency to think of self-awareness and phenomenal character as two separate phenomena: "What makes an experience conscious is not just its phenomenal quality, but also a certain *reflexive character*: a conscious mental state includes a certain inner awareness of itself" (1989, p. 98). My object in this chapter is to marshal considerations in favour of the view that phenomenal character and self-awareness are intimately linked; in fact, I want to suggest that the phenomenal character of experience is *constituted by* an implicit self-awareness of the kind explicated in Chapter Six. This chapter, in short, is a defence of the suggestion that

> . . .all subjective experience is self-conscious in the weak sense that there is something it is like for the subject to have that experience. This involves a sense that the experience is the subject's experience, that it happens to her, occurs in her stream.
>
> (Flanagan 1992, p. 194; cf. McGinn 1991b, pp. 29–30)

Here Flanagan assumes, without argument, that the what-it-is-likeness or phenomenal character of our experience is a type of self-awareness. My aim is to elaborate on this suggestion, to defend the thesis that phenomenal character is constituted by a type of self-awareness, i.e., by a low-level or implicit self-awareness that is built into every conscious state.[1]

---

1. In advancing this thesis I do not thereby presume that phenomenal character cannot be explained reductively. My thesis is that phenomenal character is constituted by – and hence explicable in terms of – implicit self-awareness, but the truth of this claim does not preclude

8.2 Phenomenal character as implicit self-awareness

I will refer to the thesis that phenomenal character is explicable in terms of implicit self-awareness as the Implicit Self-Awareness Thesis, or ISA. According to ISA, to say that there is something it is like for a creature to be in a certain mental state, is to say that the creature is implicitly aware of herself being in that mental state. The main argument for ISA is rather easily stateable, especially given the ground-clearing done; and, as I will show, a powerful consideration in favour of the thesis is that it is effectively and plausibly (though perhaps less definitively) extendible to non-perceptual conscious states, such as sensations. I begin by applying ISA to perceptual experience, after which I make some remarks on the phenomenality of conscious thought and imagination. I then apply the thesis to pains, itches, tickles, and the like.

Among those who inquire into the nature of the phenomenal character it is often supposed that the properties of perceived objects somehow *constitute* the phenomenal character of perceptual experience. According to Levine, for example, a subject's visual experience of a red diskette case has a "'reddish' character" (2001, p. 6); according to Kriegel, when a subject looks at a blue sky "there is a bluish way it is like" for the subject to have her experience (2005, p. 23); and according to Thompson, "terms such as 'greenish' and 'reddish' characterize the phenomenal character of experiences typically had by actual perceivers in response to green and red things, respectively" (2007, p. 331n7). The strategy here, as I understand it, is to explain phenomenal character in terms of our perceptual experiences somehow assuming or "taking on" the properties of the objects of experience, so that when one looks at a red diskette case the phenomenal character of one's visual experience is "reddish," whereas when one turns to look at a blue sky, the phenomenal character of one's visual experience becomes "bluish." On this view, "phenomenal character *involves* the surface qualities of which the subject of the visual experience is directly aware – these qualities at least partly *constitute* phenomenal character" (Tye 2000, p. 48).

But difficulties abound here. What it is like to see something can be neither reddish nor bluish. The diskette case and the sky are, respectively, reddish and bluish, but what it is like to see them is not. So it cannot be the case that when one looks at the diskette case, the phenomenal character of one's visual experience is

---

the possibility that implicit self-awareness is itself explicable (ultimately) in terms of something else (e.g., in terms of neural goings-on), in which case phenomenal character would be explicable (ultimately) in terms of something else. Stated otherwise, if my thesis is correct and if implicit self-awareness can be explained reductively, then, *ipso facto*, phenomenal character can be explained reductively.

reddish, whereas when one turns to look at the blue sky, the phenomenal character of one's visual experience becomes bluish.

Now those who espouse the view that visual experiences can instantiate properties like reddishness, bluishness, etc. would most likely react with indignation to this claim, arguing that they do not advocate the view that the phenomenal character of one's visual experience of a red diskette case instantiates the very same extra-mental property, reddishness, that is instantiated in the diskette case. They would insist that these are distinct properties, and that it is perfectly acceptable to say that the phenomenal character of one's visual experiences can assume properties like reddishness, bluishness, etc. as long as we are speaking in terms of reddish or bluish *feels* or reddish or bluish *sensations*.[2] In looking at the red diskette case, one has a reddish visual sensation, whereas when one turns to look at the blue sky, one has a bluish visual sensation. Thus, one's visual experience of the red diskette case just *feels* different from one's visual experience of the blue sky: one's visual experience of the diskette case feels reddish; whereas one's visual experience of the blue sky feels bluish.

No matter how it is formulated, however, this view has a number of unattractive features, the most conspicuous of which is that it is wildly at variance with visual phenomenology. Phenomenologically speaking, there is no sensation whatever involved in visual experience, except, perhaps, if one looks directly at an intense light and experiences pain in one's eyes. To be sure, the use of the expression "visual sensation" is ubiquitous in the consciousness literature, and philosophers routinely talk of having a "sensation of blueness" or a "blue sensation." But, as Merleau-Ponty has pointed out, nothing could be more confused than saying "I have a sensation of redness, of blueness" – colours are not sensations; they are things that are sensed (1962, pp. 3, 4).[3] There is, of course, something it is like to look at coloured objects, since there is something it is like to look at objects, and the objects we see come coloured.[4] But this what-it-is-likeness does not consist in the having of colour

2. See, e.g., Carruthers (2000, p. 93), Chalmers (1996, p. 5), Lycan (1996, p. 70), Tye (1995, p. 103), and Nelkin (1993, p. 420).

3. I here assume an objectivist account of colour. Various forms of irrealism, relativism, subjectivism, and dispositionalism about colour have always been popular (Galileo, Descartes, the Empiricists, etc. were proponents of such views), and this tendency shows no signs of abating: see, e.g., Hardin 1988, Boghossian and Velleman 1989, Johnston 1992, Maund 1995, and McGinn 1996. Any kind of responsible discussion of the nature of colour is beyond the scope of this book, so I will simply state my position and proceed accordingly. The position is not, as is often supposed, "anti-scientific," as Tye has ably shown (2000, ch. 7). And see Hacker 1987 and Stroud 2000 for compelling reasons for supposing that colour objectivism is true.

4. For the most part. Clear water, glass, and other transparent objects are not coloured.

this connotes the having of, or the undergoing of, discrete mental episodes or feels that correspond to the discrete colours of the objects of visual consciousness. When I see a rose, I do not have a sensation of redness or a reddish visual feel; rather, I simply see the rose, sans any sort of visual sensation.

In addition to its phenomenological implausibility, the claim that there are visual sensations falls afoul of the conventions governing the use of our sensation terms. If we dispense with the ordinary meaning of "sensation" ("feel," etc.), then perhaps it would not be inappropriate to say that there are visual sensations. However, I can find no proviso in the writings of those who currently use phrases like "visual sensation" that they are using "sensation" in a special technical sense, so one may presume that they think of themselves as deploying the concept as ordinarily conceived. But if "sensation" has its customary meaning, then no sense attaches to the claim that we have sensations of blue, of green, of lavender; for sensations always have a particular location *in* one's body, whereas simple seeings, smellings, hearings, etc. don't. Harman is admirably clear on this:

> Normally we use the term "sensation" for bodily feelings. Usually sensations have more or less definite location in one's body – a headache, a pain in one's foot, butterflies in the stomach, etc. . .. But the perception of color does not normally involve sensations in any ordinary sense of the term "sensation". When someone literally has visual sensations, they are pains or other feelings in the eye, resulting from overly bright scenes, perhaps, or by itching from allergies or minor eye injuries. Color perception does not normally involve such sensations. On seeing what appears to be a ripe tomato, one does not feel a sensation of red in one's eyes, nor is there literally a sensation or feeling at the location at which the tomato looks red.
>
> (1996, p. 7–8)[5]

One might try to avert this argument by maintaining that although our perceptual experiences typically lack sensations or feels, they nevertheless instantiate

5. Bennett and Hacker are equally explicit: "To see a red apple," they aver, "is not to have any sensations, nor is to hear a sound or to smell a smell. One has and feels sensations *in* a part of one's body, but one normally feels no sensations *in* one's eyes when one sees something, in one's ears when one hears something, or in one's nose when one smells something. There *can* be sensations in perceptual organs. One's eyes may itch, and one's ears may ache – but these sensations are not produced by the exercise of the perceptual faculty. There *can* be sensations in organs of perception produced by their exercise. Looking into a blinding light produces a sensation of being dazzled and sometimes makes one's eyes *hurt*; hearing a very loud noise produces a sensation of being deafened. But far from these sensations being constitutive of perceiving, they are accompaniments of, and impediments to, perceiving" (2003, p. 134).

properties that are *analogous to* the various extra-mental properties of the objects of perceptual consciousness. This is how a number of philosophers seem to construe particular phenomenal qualities. Peacocke (1983), for example, refers to the phenomenal character of an experience of red as "red′" ("red-prime"); and Loar (2003) refers to it as "red*". According to this view, my visual experience of a red and blue ball instantiates mental analogues of red, blue, and roundness, so that what it is like for me to see the ball, while lacking any sensation or feel, is reddish, bluish, and roundish.

I have two comments in response to this proposal. First, I would bluntly deny, again on phenomenological grounds, that our perceptual experiences instantiate such analogous properties. I, at least, do not find such properties in my perceptual experience, which should show up if they are, as they are purported to be, *experienced* aspects of conscious perceptions. Second, the idea that we can think of phenomenal character on the model of properties that are analogous to the properties – colours, contours, etc. – of the objects of consciousness is a clear instance of a fallacy bestowed upon us by the Empiricists: the fallacy, namely, of making "perception out of things perceived" (Merleau-Ponty, 1962, p. 5). I cannot give a more illuminating commentary of this fallacy than that provided by Merleau-Ponty. He writes:

> The alleged self-evidence of sensation is not based on testimony of consciousness, but on widely held prejudice. We think we know perfectly well what 'seeing', 'hearing', 'feeling' are, because perception has long provided us with objects which are coloured or which emit sounds. When we try to analyze it, we transpose these objects into consciousness. We commit. . .'the experience error', which means that what we know to be in things themselves we immediately take as being in our consciousness of them. We make perception out of things perceived.
>
> (1962, p. 5)

The fallacy of making perception out of things perceived, or "the experience error" as Merleau-Ponty calls it,[6] involves transposing the properties of objects – e.g., redness, blueness, roundness, etc. – into consciousness. And the supposition that the phenomenal character of perceptual experience can be reddish, bluish, roundish, etc. is part and parcel of this fallacy. In saying that visual experiences can instantiate properties like reddishness, one subtly converts an extra-mental property into a property of one's visual experience.

Of course, none of this is meant to suggest that conscious perceptual states lack phenomenal character. An ordinary perceptual state such as looking at a rose does

6. Place (1956) calls it the "phenomenological fallacy."

not have an isolable qualitative feel, nor does it instantiate a property analogous to redness; but there is something it is like to look at a rose. Now my basic claim is that this what-it-is-likeness is constituted by the reflexive nature of consciousness. Recall that, according to the "what-it-is-like" criterion for state consciousness proposed in Chapter Three, a mental state M of a subject S is conscious at *t* just in case there is something it is like for S to be in M at *t*. And, according to the account of state consciousness defended in Chapter Six, a mental state M of a subject S is conscious at *t* just in case S is implicitly aware of being in M at *t*. It follows (by simple biconditional introduction) that:

> (ISA_1) There is something it is like for one to be in a mental state M at *t* if and only if one is implicitly aware of oneself being in M at *t*.

According to this thesis, then, a subject is in a phenomenal perceptual state simply in virtue of being in a reflexive perceptual state.

ISA_1, however, is silent on the *nature* of phenomenal character. It says (merely) that the phenomenal character and the implicit self-awareness of a conscious experience necessarily co-occur, or that one (causally?) depends on the other. It does not answer the question of what phenomenal character *is*; nor does it answer the question of *what* it is like for a subject to have a particular perceptual experience. Can these questions be answered within the framework of the Brentanian-Sartrean reflexive account of state consciousness sketched in Chapter Six? I think they can, and rather neatly in fact. Phenomenal character, I want to suggest, *just is* – i.e., is constituted by – the implicit self-awareness involved in conscious states. And *what* it is like for a subject to perceive *x* is itself constituted by the reflexive nature of consciousness, in the following sense: if a visual experience – e.g., of a chair – is like something to me, then what it is like for me to see the chair is like *seeing the chair*. Using Brentanian terminology, this proposal may be formulated as follows:

> (ISA_2) In consciously seeing *x* (hearing *x*, etc.), one is primarily conscious of *x* and secondarily conscious of oneself seeing *x* – one's visual consciousness of *x* is *like* something, i.e., like seeing *x*.

And using Sartrean terminology, we get:

> (ISA_3) When one consciously sees *x* (hears *x*, etc.), one has a positional or thetic consciousness of *x* – *x* is what S attends to – and a non-positional or non-thetic consciousness of oneself seeing *x* – one's seeing *x* is *like* something, i.e., like seeing *x*.

That seeing *x* is like seeing *x* might strike some as trivial (or even tautologous), but it's a natural consequence of the assumption that conscious states are reflexive. If a subject consciously sees *x*, the subject "lives through" the visual experience and so

is immediately experientially acquainted with it. Put differently, the subject's visual experience is pre-reflectively *given* to the subject. But insofar as the subject's visual experience is given to the subject, it is given to the subject as a particular perceptual experience – in this case an experience of seeing *x* – from which it follows that seeing *x* is what it is like for the subject to see *x*.

ISA is, in a sense, a deflationary account of phenomenal character, but it is neither trivial nor uninformative. It tells us what phenomenal character is, i.e., it tells us that phenomenal character is constituted by the implicit self-awareness involved in conscious states; and it tells us what it is like for a subject to be conscious of something, i.e., it tells us that what it is like for a subject to consciously see *x* (consciously hear *x*, etc.) is like seeing *x* (hearing *x*, etc.). That ISA is uninformative might be levelled against this latter component of the thesis in particular, but such a charge is question-begging, since it presupposes that, as regards what it is like for a subject to perceive something, more information is in the offing than this part of ISA provides. But the whole point of this part of the thesis is precisely that what it is like for a subject to perceive something is itself constituted by – and therefore "bottoms out" in – the reflexive structure of consciousness.

As I have already suggested, one of the benefits of this approach to the analysis of phenomenal character is that it provides an intuitively appealing explanation of the what-it-is-likeness of conscious perceptual states without postulating intuitively unappealing intra-mental "feels" or qualia that correspond to the objects, or the properties of the objects, in one's perceptual field. Consider Nelkin's claim that "looking at the hands of [a] dial watch just is qualitatively different from looking at a digital watch face. . . . Such experiences just 'feel' different from each other" (1993, p. 420). The view expressed by this claim answers to nothing in our experience: what it is like to look at the hands of a dial watch does not *feel* like anything, and the same is true of looking at a digital watch face. Yet looking at these objects is, indeed, like something – when I look at a digital watch face (the hands of a dial watch, etc.) my visual experience has phenomenal character. According to ISA, my looking at the digital watch face is like something insofar as I am implicitly aware of looking at the digital watch face; and what it is like for me to look at the digital watch face is like looking at the digital watch face.

Another attraction of ISA is that it affords a relatively straightforward explanation of the phenomenality of conscious thinking (including introspection, since introspection is a kind of conscious thinking, a matter of focusing one's attention on something;[7] only, in introspection, the object of one's conscious thinking is

7. And not, for example, a form of perception or inner vision (supporters of this kind of view of introspection, which is thought to have originated with Locke (1975), include

oneself, or one's own occurrent or non-occurrent mental state – one's fear, happiness, pain, etc. – not an extra-mental object). If one embraces, as Woodruff Smith does for example, the view that each thought has its own distinctive cogitative feel or quale, then one is committed to affirming that thinking that Russell was subtler than Wittgenstein feels different from thinking that Wittgenstein was subtler than Russell (see Woodruff Smith 1989, p. 96). But one can deny that these thoughts have distinctive feels without denying that conscious thinking lacks phenomenality. On the proposed view, there is something it is like to consciously think about *x* insofar as one is implicitly aware of oneself thinking about *x*. The phenomenal character of the experience consists in the subject of the act of thinking being implicitly aware of herself thinking about something (be it an extra-mental object or, as in introspection, one's own mental state).

The same, I suggest, is true of imagination. There is something it is like to imagine a pink elephant, and there is something it is like to imagine a blue elephant, but what it is like to imagine a pink elephant does not feel different from what it is like to imagine a blue elephant; indeed, imagining a pink elephant does not *feel* like anything, and the same is true of imagining a blue elephant. The what-it-is-likeness of these acts of imagination consists in the subject being implicitly aware of herself imagining something.

By now it should be clear that adopting ISA can lead to important gains in theoretical parsimony. By way of illustration, I should like to briefly contrast ISA with Woodruff Smith's account of phenomenal character. According to Woodruff Smith, all conscious states have phenomenal qualities or qualia: "every sensation, perception, desire, or thought, insofar as it is conscious, has a certain quality of 'what it is like', or what it 'feels' like, to have that type of experience" (1989, p. 95). Now Woodruff Smith holds, additionally, that every conscious state also has a certain reflexive character; thus, on his view, all conscious states have both a phenomenal quality and a reflexive character, each of which is essential to consciousness.[8]

Armstrong (1968, 1981) and Lycan (1996)). Introspection involves no perceptual organ (as, for example, vision does); there is no inner organ, analogous to one's eyes, that one uses to attend to one's own mental states (as one uses one's eyes to see things). So, too, introspection does not involve observation: in attending to one's occurrent feeling of happiness, for example, one does not observe it. Critics of the perceptual model include Bennett and Hacker (2003, §3.6), Moran (2001, §1.2), Searle (1992, pp. 143–4), and Shoemaker (e.g., 1996b).

8. In the course of expounding the HOT theory, Rosenthal similarly invokes a distinction between self-awareness and phenomenality. On Rosenthal's account, conscious states have qualitative or phenomenal properties of which subjects can become aware (see Rosenthal 1986, pp. 351–2; 2002c, pp. 412–14; cf. Gennaro 1996, ch. 6).

The phenomenal quality of seeing a yellow rose, Woodruff Smith says, answers to the following phenomenological description:

(1) Phenomenally (I see this yellow rose),

while the reflexive character of the same experience answers to the following phenomenological description:

(2) In this very experience I see this yellow rose.

According to ISA, (1) and (2) express the same proposition. To say that, phenomenally, I see this yellow rose is to say that, in this very experience, I see this yellow rose. Woodruff Smith, by contrast, holds that phenomenality is something very different from reflexivity, which means that, for him, (1) and (2) express different propositions. For Woodruff Smith, (1), unlike (2), expresses a fact about an intramental property: in seeing a yellow rose, one has a distinctive feel or quale. ISA, on the other hand, expels qualia (feels, etc.) from consciousness,[9] and explains the what-it-is-likeness of conscious perceptual states in terms of their reflexive character, i.e., in terms of creatures being implicitly aware of being conscious of things.

8.2.1 Differences in phenomenal character

A largely unchallenged assumption in discussions of phenomenal character is that, in addition to *course-grained* or *inter-modal* differences in phenomenal character (e.g., the difference between what it is like to *see* an apple and what it is like to *feel* an apple), there are *fine-grained* or *intra-modal* differences in phenomenal character (e.g., the difference between what it is like to see a *red apple* and what it is like to see a *green apple*, or the difference between what it is like to see an apple at *t* and what it is like to see the same apple at *t+*, where, between *t* and *t+*, the apple has undergone a perceptible surface change of some sort).[10] In fact, this assumption strikes many as virtually platitudinous. Envisaging no uncertainties or objections, philosophers who embrace it simply *announce* that there are inter-modal and intra-modal differences in phenomenal character. Here, for example, is Langsam on inter-modal differences:

---

9. Cf. Sartre: "[t]he first procedure of a philosophy ought to be to expel things from consciousness and to reestablish its true connection with the world, to know that consciousness is a positional consciousness of the world" (1956, p. li).

10. Many extend this view to purely cognitive states, arguing that what it is like to think that *p* is different from what it is like to think that *q* (e.g., Woodruff Smith 1989, Strawson 1994, Chalmers 1996, McCulloch 1999, Langsam 2000, Pitt 2004, Zahavi 2005).

> What is obvious about perceptual experiences is that they are very different from thoughts. The differences with which I am concerned are the differences in their phenomenal character: what it is like to have a perceptual experience is very different, is obviously different, from what it is like to have a thought.
>
> (2000, p. 273)

And here is Byrne on intra-modal differences:

> Imagine the experiences of looking, on separate occasions, at three coloured chips on a neutral background. The first two chips are blue, and the third red, and the subject has normal vision. All can agree that, as far as "what it's like" to undergo the experiences is concerned, seeing the first chip and seeing the second having something in common that seeing the third lacks.[11]
>
> (Byrne 2001, p. 206)

The assumption that there are both inter-modal and intra-modal differences in phenomenal character seems eminently plausible, and any satisfactory theory of phenomenal character should have the resources to account for these differences. But there is another important fact about the phenomenal character of perceptual experience that any such theory should be able to explain, namely, that concept acquisition can affect differences in phenomenal character. Visual states (auditory states, etc.) are frequently concept-imbued, in the sense that we see things *as* being so-and-so (see *that* things are so-and-so, etc.). We see the piece of electronic equipment on the table *as* a computer; we see *that* the bird in the distance is an osprey; and so on. In view of this, it seems reasonable to suppose that what it is like for one to see an object upon which one has brought a concept to bear can differ from what it is like for one to see an object upon which one has not brought a concept to bear. Similarly, it is plausible that what it is like for one to see *x* (hear *x*, etc.) prior to one's conceptualization of *x* can differ from what it is like for one to see *x* after one's conceptualization of *x*. Consider, in this connection, the following anecdote from Carruthers about his first job in the wilds of Scotland, where there was little to do but bird-watch:

---

11. Tye makes a similar claim regarding intra-modal differences in phenomenal character: "As you view the scene before your eyes and how things look to you, necessarily, if any of the qualities of which you are *directly* aware change, then the phenomenal character of your experience changes. Consider, for example, the facing surface of a ripe tomato in a bowl before you. In attending to the color of the tomato, you are directly aware of a certain quality, *Q*, as covering that surface. You experience each just noticeable part of the surface as having Q. . . . But change *Q* – for example, by changing the color of the tomato or by donning color-inverting lenses – and what it is like for you in viewing the tomato necessarily changes" (2000, p. 48).

> At first I saw crowds of little grey birds on the beach, but I later came to *see* the beach as populated by plovers, knots, dunlins and red-shanks. As a result of concept-learning, the differences between the birds came to leap out at me in a phenomenologically salient way; I *saw* them as distinct. It soon became barely intelligible to me how I could ever have confused a plover with a dunlin, they looked so different. (2000, p. 131)

What it was like for Carruthers to see a dunlin at *t* differed from what it was like for him to see a dunlin at *t*+, since at *t* he had no concept of a dunlin, whereas at *t*+ he did. At *t*, he saw the dunlin simply as a nondescript grey bird, whereas at *t*+ he saw it as a *dunlin*. It seems, then, that modifications in one's conceptual repertoire can yield intra-modal differences in the phenomenal character of perceptual experience.[12]

But how, one might justifiably wonder, can ISA explain this fact? And how can it explain inter-modal and intra-modal differences in phenomenal character? After all, if, for example, seeing a blue chip does not, as ISA says, *feel* different from seeing a red chip – i.e., if these experiences do not have distinct sensory feels or qualia – then what, according to ISA, does the intra-modal difference in phenomenal character between these experiences consist in?

As regards both inter-modal and intra-modal differences in phenomenal character, this issue need not detain us unduly; for ISA states that phenomenal character is constituted by implicit self-awareness, and so wherever there is a difference in implicit self-awareness, there will be a difference in phenomenal character. And differences in implicit self-awareness issue from either differences in perceptual modality or differences in the object (or properties) being perceived. For example, according to ISA, the phenomenal character of one's *visual* experience of a blue chip can be explained in terms of one's implicit awareness of oneself *seeing* (or looking at) a blue chip, while the phenomenal character of one's *tactile* experience of a blue chip can be explained in terms of one's implicit awareness of oneself *touching* (or feeling) a blue chip. So the inter-modal difference in phenomenal character between these experiences issues from the different perceptual modalities in question: visual and tactile. One has distinct perceptual experiences by way of distinct perceptual modalities; and since one is implicitly aware of – i.e., lives through – these distinct experiences, and since phenomenal character is constituted by implicit self-awareness, a difference in phenomenal character will issue from the difference in implicit self-awareness, which, in turn, derives from a difference in perceptual modality.

---

12. Rosenthal has also pressed this claim (e.g., 1986, p. 350; 2002c, §IV).

Roughly parallel considerations apply to intra-modal differences in phenomenal character. According to ISA, the phenomenal character of one's visual experience of a blue chip can be explained in terms of one's implicit awareness of oneself seeing *a blue chip*, whereas the phenomenal character of one's visual experience of a red chip can be explained in terms of one's implicit awareness of oneself seeing *a red chip*. The difference in phenomenal character between these two visual experiences issues from the objects being perceived. Thus, ISA is entirely consistent with the assumption that the phenomenal character of an experience e in sense modality m differs from the phenomenal character of an experience e^* in sense modality m insofar as the object of e differs from the object of e^*. However, ISA does not construe the difference in phenomenal character between e and e^* in terms of their having different feels or qualia. Rather, ISA construes the difference in phenomenal character between e and e^* in terms of e and e^* having different objects (or properties).[13]

ISA can also readily explain how it is that concept acquisition can affect changes in phenomenal character. If a subject perceives an object x upon which he has brought no concept to bear, then the subject will be implicitly aware of his nonconceptual perceptual experience of x: it will be immediately given to him. If the subject subsequently perceives x after having subsumed x under a concept, then he may have a concept-imbued perceptual experience of x (e.g., he may see x *as* such-and-such), and hence will be implicitly aware of a concept imbued perceptual experience of x. Since there is a difference in implicit self-awareness between the subject's nonconceptual experience of x and his conceptual experience of x, there will also be a difference in phenomenal character between the two experiences. Thus, the difference in phenomenal character derives from a difference in implicit self-awareness, which, in turn, derives from different experiences: a conceptual experience of x and a nonconceptual experience of x.

8.2.2 Pains and other bodily sensations

I stated in the preamble that ISA is plausibly extendible to all types of conscious states, and so in this section I extend the analysis to the phenomenal character of bodily sensations, such as pains, itches, tickles, and the like.

---

13. I am here using the term "object" rather loosely to mean, roughly, "content." If I hallucinate a pink elephant at t and a blue elephant at $t+$, the object of my successive visual experiences does not differ, since neither experience *has* an object. Nevertheless, both of my hallucinatory experiences have content, and the content – and hence the phenomenal character – of my experience at t is different from the content of my experience at $t+$.

A provisional statement of ISA vis-à-vis pain might proceed as follows. When I am in a conscious pain state, I am conscious of some bodily state or condition, usually an injured, strained, or diseased bodily state or condition. But I am also implicitly aware of my pain, i.e., my pain is like something. Now I can reflect on or think about my pain, pondering the peculiarities of the somatic thumping of my headache, in which case I am introspectively aware of my pain. But even in such cases, there is an implicit self-awareness that accompanies my conscious state: I am implicitly aware of my act of introspection, i.e., my act of introspection is like something.

Now an objection that might be raised here is that this sort of analysis does not account for the phenomenal character of pains, itches, tickles, and the like because pains and other bodily sensations are different in that, unlike perceptual states and conscious thoughts, they possess what might be called "raw feels." So even if one were implicitly aware of oneself being in pain, this would not account for the *painfulness* of pain, i.e., its raw feel.

It is by no means evident, however, why the implicit self-awareness of pain need be (could be) anything other than the what-it-is-likeness of the pain (and vice versa). The painfulness of the pain *is* an implicit awareness of the pain. But where the raw feel? Why a raw feel, except to introduce an unnecessary intra-mental property? When I am in a conscious pain state, I am conscious of some bodily region *r*; and in being conscious of *r*, I am implicitly aware of myself as being conscious of *r* – my being conscious of *r* is like something.[14] There is no difference here between pain and ordinary perception or imagination.

Of course, ordinarily, when one is in pain, one immediately notices the pain – which is unsurprising given its indisputable unpleasantness – and goes into "introspective mode." In this way, pain is different from ordinary perception and imagination in that it is much less self-effacing. When one imagines *x*, visually perceives *x*, auditorily perceives *x*, etc. one rarely notices oneself imagining *x*, seeing *x*, hearing *x*, etc., but when one has a pain in *x*, one typically shifts one's attention from *x* to one's consciousness of *x*, i.e., to one's pain. But this fact about pain does not undermine the view being advanced here. In ordinary perception, as I remarked, it is not difficult to shift one's attention from the object of one's act of looking, say, to one's act of looking itself. Normally we do not do this because it is the nature of a perceptual state to recede into the background just to the extent that the subject is caught up, via that perceptual state, with some object, event, or state of affairs in the world. So even though we are less prone, in ordinary perception, to attend to our perceptual states themselves, we certainly *can* attend to them.

---

14. I owe this rather colourful way of stating the point to Phil Dwyer (in conversation).

There is a plausible just-so story to be offered about the evolutionary benefits of the self-effacing nature of ordinary perception and the non-self-effacing nature of pain. There are definite survival advantages in being aware of bodily disturbances, and in being able to distinguish disturbances that are potentially fatal from those that are not. Pain is nature's way of signifying these disturbances, which is why pain, as Dennett (1991, p. 61) has put it, has been such an "evolutionary boon." Creatures with congenital insensitivity to pain rarely survive into adulthood. But notice that creatures typically want to rid themselves of their pain states for the obvious reason that pain is unpleasant. If pain were not unpleasant, then a creature would have no reason to combat it. But pain's unpleasantness is just its way of drawing attention to itself. And if a conscious state draws attention to itself in an unpleasant way, then the creature whose conscious state it is will take steps to eliminate that state.[15] Thus, since pain states draw attention to themselves in an unpleasant way, the creature whose pain it is will take action to assuage the pain, perhaps by attending to the bodily injury causing the pain. And this, in turn, will enhance the creature's chances of survival.[16]

---

15. Cf. Bennett and Hacker: "If a sensation is intense, it cannot but catch one's attention. The more intense it is, the more it thrusts itself upon our attention. If it is momentary, like a sudden twinge, one cannot but be aware of it. If it is persistent, one cannot but be conscious of it. The more intense the sensation, the more difficult it is to attend to anything else or, indeed, to think of anything else. It is no coincidence that we speak of being 'in the grip' of pain. The less intense a sensation, the easier it is for us to have our attention *distracted* from it or to *forget* it" (2003, p. 259).

16. A philosophical bugbear lurking in the vicinity here is worth commenting on. It has been claimed that we should be concerned about whether conscious states have *any* evolutionary value (e.g., Chalmers 1996, pp. 4–5), and so one might object that I am moving forward too quickly, that prior to offering a story about the evolutionary value of particular sensory or perceptual states, one must first offer a plausible story about the evolutionary value of conscious states generally. The question of whether conscious states confer a survival advantage on creatures that have them issues, at least in part, from the stultifying spectre of the alleged logical possibility of zombies. The idea is that there is room to doubt that conscious states confer a survival advantage because there is nothing logically awry about supposing that it is possible for a creature of exhibit conscious behaviour, i.e., to behave the same way we do, in the complete absence of consciousness. I do not take seriously the claim that zombies are logically possible, but even if they are, such a weak conceivability consideration is no reason to suppose that conscious states do not have a biological "function" (just as the alleged logical possibility of our being envatted brains is no reason to suppose that there is no "external" world). There are, of course, biological questions about how, for example, primitive organisms evolved into more sophisticated organisms that can be said to have conscious states, but no metaphysical mysteries attach to these sorts of questions. Moreover, evolution is "opportunistic": it goes with what works. So even if, to use Dretske's words, "something else besides experience would enable us to do some things. . .this would not show

Ordinary perceptual states, by contrast, are not unpleasant, which means that the subject will have no reason to combat them. But then the survival benefits of having such perceptual states is fairly obvious. If a gazelle takes evasive action upon espying an attacking lion, it seems perfectly appropriate to explain the gazelle's behaviour by reference to the gazelle's *seeing* the lion – and thus by reference to a particular perceptual state. This suggests that ordinary perceptual states are self-effacing for good evolutionary reasons: they recede to the background, make themselves unobtrusive, just to the extent that this takes the subject transparently to the perceptual state's content, i.e., to the worldly item at which the perceptual state is directed, which, in turn, enhances the subject's chances of survival.

There is, then, a story to be told about how pain differs from ordinary perception, but this story is not one about how pains, unlike ordinary perceptual states, have raw feels. A pain state does not have a raw feel, where the having of a raw feel connotes the having of some kind of intra-mental property that outstrips the implicit self-awareness involved in that state.

I suspect that any lingering reluctance to countenance this view derives from the firmly entrenched, yet no less mistaken view that pains are mental objects of some kind. We regularly employ expressions like "I feel a dull, throbbing pain," i.e., expressions in which the word "pain" is grammatically the accusative object. But, as I have already mentioned, we should not be misled by the grammar of these expressions.[17] To *feel* a pain is simply to be in a conscious pain state, and to be in a conscious pain state is simply to be conscious of some bodily region. Of course, when we reflect on or think about our pains, it is not inappropriate to refer to the pain as the object of one's thought. But just as our ability to think about our perceptual states is no reason to suppose that perceptual states are mental objects, so our ability to think about our pain states is no reason to suppose that pain states are mental objects.

8.3 Some brief remarks on privacy

Since ISA obviates the need to explain phenomenal character by postulating intra-mental properties like raw feels or qualia (I will hereafter subsume "raw

---

that experience didn't have a function. All it would show is that there was more than one way to skin a cat – more than one way to get the job done" (1997, p. 14). For more on the evolutionary advantage of consciousness, see Bennett and Hacker 2003, §11.4; and Tye 1996.

17. See Douglas 1998 for an argument purporting to show that, in many cases, philosophers are led to the view that pains are mental objects by the misleading surface structure of the grammar of sensation sentences.

feels" under the rubric "qualia"), it also has the decided advantage of avoiding the seemingly interminable disputes over whether qualia are essentially private. The whole issue of privacy regarding qualia arises because qualia are construed as intra-mental properties to which subjects stand in a certain sort of (as yet undetermined) relation,[18] but the phenomenal character or what-it-is-likeness of a conscious experience, according to ISA, is not part of a realm of subjective, intra-mental properties to which subjects stand in a relation. In fact, the phenomenal character of a conscious experience is not something the subject is aware of at all in having the experience. In seeing *x* (hearing *x*, having a pain in *x*, etc.), the subject is implicitly aware of herself seeing *x* (and so the subject might be said to stand in a relation to herself), but the subject is not – in addition to being implicitly aware of herself seeing *x* – aware, somehow, of being implicitly aware of herself seeing *x*. The phenomenal character of a perceptual or non-perceptual experience, on this view, is that in virtue of which we perceptually or non-perceptually experience a worldly (or mental) item in a certain way; it is not something of which we are aware in having the experience.

Now it might be protested that I have ignored the sense of privacy at issue in current debates, a sense of privacy that poses a problem even for those who reject raw feels, qualia, and the like. This sense of privacy is an artefact of the assumption that there is a pronounced asymmetry between first-person and third-person apprehension of conscious states, such that conscious states are essentially *subjective*. This sense of subjectivity is usually stated in terms of subjects having *privileged access* to their own conscious states. Searle puts it this way:

> My present state of consciousness is a feature of my brain, but its conscious aspects are accessible to me in a way that they are not accessible to you. And your present state of consciousness is a feature of your brain and its conscious aspects are accessible to you in a way that they are not accessible to me.
>
> (1984, p. 25)

Putting aside the dubious claim that conscious states are features of brains,[19] however, it makes scant sense to say conscious states are specially accessible to their

---

18. For an especially vivid portrayal of a quale as an inner mental property to which the subject stands in a relation, consider Lycan's definition of a quale in the strict sense: "A quale is the introspectible monadic qualitative property of what seems to be a phenomenal individual, such as the color of what Russell called a visual sense datum. For example, if S is visually healthy and looking at a ripe tomato in good light, the tomato will look red to S, and if S focuses her introspective attention on the corresponding subregion of her visual field, S will see that subregion as an individual red patch having a roundish shape. The redness of that phenomenal patch is the quale of the containing visual sensation" (1996, pp. 69–70).

19. Conscious states are states of whole animals, of creatures, not states of brains; and, to adopt Searle's unusual wording, consciousness is a feature of creatures, not a feature of brains.

subjects. Of course, only *I* can have my conscious states; and, as I have argued, it is true that a subject is conscious of *x* just in case the subject *knows* she is conscious of *x*. But the fact that only I can have my conscious states is an unpuzzling conceptual truth,[20] and the fact that a subject is conscious of *x* just in case the subject knows she is conscious of *x* is no reason to conclude that conscious states are specially accessible to their subjects.

Perhaps the idea is that subjects have privileged access to their own conscious states in the sense that the subject can know his conscious states non-inferentially, while others cannot. But this is unsupported dogma. If I see a person writhing on the ground, moaning, and clutching his chest, I do not infer from his behaviour that he is experiencing chest pain; rather, I simply *see* that he is experiencing chest pain.[21] It is true that people can suppress their emotions and feel pain and other bodily sensations and not show it; and it is also true that sometimes we have to infer that someone is in a particular conscious state. But it does not follow that subjects have privileged access to their own conscious states, where this implies that there are certain inner aspects of experience that, *as a matter of logic*, are permanently inaccessible to outside observers.

---

20. Though it does not follow that two or more subjects cannot have the *same* conscious state. Following Wittgenstein (1953, §253), Malcolm has persuasively argued that "there is no sense of the expression 'same pain' such that it is impossible for two people to have the same pain" (1977, pp. 121–2). In addition, it is generally thought that if subjects S and S_1 both claim to have, say, a dull, throbbing headache, then S and S_1 have qualitatively identical pains that are numerically distinct. But this way of looking at the matter is, arguably, mistaken. Bennett and Hacker, for example, have argued that "the distinction between qualitative and numerical identity, which applies to material objects, does not apply to properties, such as colours, or to pains. Two material objects may be qualitatively identical but numerically distinct. Neither colours nor pains admit of this distinction. . ." (2003, p. 96n62; cf. Malcolm 1977, p. 122n32). They offer the following analogy: "Two red cherries don't have different colours just because the red of the first cherry 'belongs' to the first cherry and the red of the second 'belongs' to the second" (2003, p. 95). (Incidentally, if pains do not admit of the distinction between qualitative and numerical identity, then pains cannot be identical to neural states, since neural states *do* admit of the distinction between qualitative and numerical identity.)

21. As Wittgenstein says, choosing the example of fear: "In general I do not surmise fear in him – I *see* it. I do not feel that I am deducing the probable existence of something inside from something outside; rather it is as if the human face were in a way translucent and that I were seeing it not in reflected light but rather in its own" (1980, §170). And, more generally: "Look into someone else's face, and see the consciousness in it, and a particular *shade* of consciousness. You see on it, in it, joy, indifference, interest, excitement, torpor and so on" (1967, §220). These remarks are sometimes construed as an endorsement of behaviourism, but Wittgenstein is not endorsing behaviourism with these remarks, but stating truisms.

8.4 Conclusion

My principal aim in this book was to expound and defend two primary theses: (1) that all conscious states are reflexive, in the sense that they have a built in implicit self-awareness such that, in having a conscious experience, a subject is immediately and non-objectually acquainted with the experience as hers; and (2) that the phenomenal character of conscious experience is constituted by the implicit self-awareness involved in conscious states.

The former thesis has a decided advantage over its rivals: it respects the Transitivity Principle while avoiding those problems that arise in connection with construing self-awareness along subject-object lines. And the latter thesis – what I have called ISA – is as plausible as (if not more plausible than) any other and merits serious attention. Fodor has remarked that "[t]he form of a philosophical theory, often enough, is: *Let's try looking over here*" (1981, p. 31). I think the case I have made for ISA creates a presumption, but I will be satisfied if it provides a direction future work on the subject of phenomenal character might take. Extant accounts of phenomenal character – so I contend and have documented (and would be prepared to further document) – routinely deny phenomenologically obvious phenomena (e.g., that the what-it-is-likeness of ordinary perceptual states does not consist in the having of feels or sensations). A major virtue of ISA is that it does not do this. On the contrary, it affords us a unified account of the nature of phenomenal character that respects the appearances.

Bibliography

Antony, M. (2001). Is 'Consciousness' Ambiguous? *Journal of Consciousness Studies 8:* 19–44.

Aquila, R. (1988). The Cartesian and a Certain 'Poetic' Notion of Consciousness. *Journal of the History of Ideas 49:* 543–62.

Aquila, R. (1990). Consciousness and Higher-Order Thoughts: Two Objections. *American Philosophical Quarterly 27:* 81–7.

Armstrong, D.M. (1968). *A Materialist Theory of the Mind*. London: Routledge.

Armstrong, D.M. (1981). What is Consciousness? In Block, Flanagan, and Güzeldere (1997), pp. 721–8.

Armstrong, D.M. and N. Malcolm (1984). *Consciousness & Causality*. Oxford: Blackwell.

Baier, A. (1990). What Emotions are About. In J. Tomberlin(Ed), *Philosophical Perspectives 4:* 1–29. Atascadero: Ridgeview.

Baker, L.R. (2000). *Persons and Bodies: A Constitution View*. Cambridge: Cambridge University Press.

Balog, K. (1999). Conceivability, Possibility, and the Mind-Body Problem. *Philosophical Review 108:* 497–528.

Bennett, J. (1966). *Kant's Analytic*. Cambridge: Cambridge University Press.

Bennett, M.R. and P.M.S. Hacker (2003). *Philosophical Foundations of Neuroscience*. Oxford: Blackwell.

Bermúdez, J. (1998). *The Paradox of Self-Consciousness*. Cambridge MA: MIT Press.

Block, N. (1983). Mental Pictures and Cognitive Science. *Philosophical Review* **92**: 499–541.

Block, N. (1994). Qualia. In S. Guttenplan(Ed.), *Blackwell Companion to the Philosophy of Mind*, pp. 514–20. Oxford: Blackwell.

Block, N. (1995). On a Confusion about a Function of Consciousness. In Block, Flanagan, and Güzeldere (1997), pp. 375–415.

Block, N. (1996). Mental Paint and Mental Latex. In E. Villanueva(Ed.), *Philosophical Issues 7:* 19–49. Atascadero: Ridgeview.

Block, N., O. Flanagan, and G. Güzeldere (Eds.) (1997). *The Nature of Consciousness: Philosophical Debates*. Cambridge MA: MIT Press.

Boghossian, P. (1990). The Status of Content. *Philosophical Review 99:* 151–84.

Boghossian, P. and J.D. Velleman (1989). Colour as a Secondary Quality. *Mind 98*: 81–103.

Bower, T.G.R. (1972). Object Perception in Infants. *Perception 1:* 15–30.

Brentano, F. (1874). *Psychology from an Empirical Standpoint*. Ed. L.L. McAlister. Trans. A.C. Rancurello, D.B. Terrell, and L.L. McAlister. London: Routledge, 1973.

Brentano, F. (1982). *Descriptive Psychology*. [Trans. Benito Müller(Ed.)]. London: Routledge.

Broad, C.D. (1925). *Mind and Its Place in Nature*. London: Routledge & Kegan Paul.

Brueckner, A. (2001). Chalmers's Conceivability Argument for Dualism. *Analysis 61*: 187–93.

Burge, T. (1997). Two Kinds of Consciousness. In Block, Flanagan, and Güzeldere (1997), pp. 427–34.

Byrne, A. (1997). Some Like it HOT: Consciousness and Higher-Order Thoughts. *Philosophical Studies 86*: 103–29.

Byrne, A. (2001). Intentionalism Defended. *Philosophical Review 110*: 199–240.

Byrne, A. (2004). What Phenomenal Consciousness is Like. In Gennaro (2004a), pp. 203–25.
Campbell, S. (2004). Seeing Objects and Surfaces, and the 'In Virtue of' Relation. *Philosophy 79*: 393–402.
Carman, T. (2005). On the Inescapability of Phenomenology. In D. Woodruff Smith and A. Thomasson (Eds.), *Phenomenology and the Philosophy of Mind*, pp. 67–89. Oxford: Clarendon.
Carruthers, P. (1989). Brute Experience. *Journal of Philosophy 89*: 258–69.
Carruthers, P. (1996). *Language, Thought and Consciousness: An Essay in Philosophical Psychology*. Cambridge: Cambridge University Press.
Carruthers, P. (2000). *Phenomenal Consciousness: A Naturalistic Theory*. Cambridge: Cambridge University Press.
Castañeda, H.-N. (1966). 'He': A Study in the Logic of Self-Consciousness. *Ratio 8*: 130–57.
Caston, V. (2002). Aristotle on Consciousness. *Mind 111:* 751–815.
Casullo, A. (1987). A Defense of Sense-Data. *Philosophy and Phenomenological Research 48*: 45–61.
Chalmers, D. (1996). *The Conscious Mind: In Search of a Fundamental Theory*. Oxford: Oxford University Press.
Chisholm, R. (1957). *Perceiving*. Ithaca: Cornell University Press.
Chisholm, R. (1981). *The First Person: An Essay on Reference and Intentionality*. Minneapolis: University of Minnesota Press.
Churchland, P.M. (1995). *The Engine of Reason, the Seat of the Soul: A Philosophical Journey Into the Brain*. Cambridge MA: MIT Press.
Clark, A. (2000). A Case Where Access Implies Qualia? *Analysis 60*: 30–8.
Collins, A. (1987). *The Nature of Mental Things*. Notre Dame: University of Notre Dame Press.
Cowley, F. (1968). *A Critique of British Empiricism*. London: Macmillan.
Cowley, F. (1991). *Metaphysical Delusion*. Buffalo: Prometheus Books.
Crane, T. (2003). The Intentional Structure of Consciousness. In Q. Smith and A. Jokic (Eds.), *Consciousness: New Philosophical Perspectives*, pp. 33–56. Oxford: Clarendon.
Crane, T. and D.H. Mellor (1990). There is No Question of Physicalism. *Mind 99*: 185–206.
Crick, F. and C. Koch (1990). Towards a Neurobiological Theory of Consciousness. *Seminars in the Neurosciences 2*: 263–75.
Dainton, B. (2000). *Stream of Consciousness: Unity and Continuity in Conscious Experience*. London: Routledge.
Damasio, A. (1994). *Descartes' Error: Emotion, Reason, and the Human Brain*. New York: G.P. Putnam's Sons.
Dennett, D. (1969). *Content and Consciousness*. London: Routledge.
Dennett, D. (1978). *Brainstorms: Philosophical Essays on Mind and Psychology*. Cambridge MA: MIT Press.
Dennett, D. (1987). *The Philosophical Lexicon*. Newark: University of Delaware.
Dennett, D. (1988). Quining Qualia. In Block, Flanagan, and Güzeldere (1997), pp. 619–42.
Dennett, D. (1991). *Consciousness Explained*. Boston: Little, Brown and Company.
Dennett, D. (1993). The Message is: There is no Medium. *Philosophy and Phenomenological Research 53:* 919–31.
Dennett, D. (1998). *Brainchildren: Essays on Designing Minds*. Cambridge MA: MIT Press.
Douglas, G. (1998). Why Pains are Not Mental Objects. *Philosophical Studies 91*: 127–48.
Dretske, F. (1969). *Seeing and Knowing*. Chicago: University of Chicago Press.

Dretske, F. (1993). Conscious Experience. *Mind 102*: 263–83.

Dretske, F. (1994). Mind and Brain. In R. Warner and T. Szubka (Eds.), *The Mind-Body Problem: A Guide to the Current Debate*, pp. 131–6. Oxford: Blackwell.

Dretske, F. (1995). *Naturalizing the Mind*. Cambridge MA: MIT Press.

Dretske, F. (1997). What Good is Consciousness? *Canadian Journal of Philosophy 27*: 1–15.

Dretske, F. (1999). The Mind's Awareness of Itself. *Philosophical Studies 95:* 103–24.

Dretske, F. (2003). Experience as Representation. In E. Sosa and E. Villaneuva (Eds.), *Philosophical Issues 13:* 67–82. Oxford: Blackwell.

Dretske, F. (2004). Change Blindness. *Philosophical Studies 120:* 1–18.

Droege, P. (2003). *Caging the Beast: A Theory of Sensory Consciousness*. Amsterdam: John Benjamins Publishing Company.

Drummond, J.J. (2006). The Case(s) of (Self-)Awareness. In Kriegel and Williford (2006), pp. 199–220.

Dwyer, P. (1990). *Sense and Subjectivity: A Study of Wittgenstein and Merleau-Ponty*. Leiden: E.J. Brill.

Edelman, G. (1989). *The Remembered Present: A Biological Theory of Consciousness*. New York: Basic Books.

Farrell, B. (1950). Experience. *Mind 59*: 170–98.

Fichte, J. (1970). *Science of Knowledge*. [Trans. P. Heath and J. Lachs]. New York: Appleton-Century-Crofts.

Field, J. (1976). Relation of Young Infants' Reaching Behaviour to Stimulus Distance Solidity. *Developmental Psychology 12:* 444–8.

Flanagan, O. (1992). *Consciousness Reconsidered*. Cambridge MA: MIT Press.

Fodor, J. (1981). *Representations: Philosophical Essays on the Foundations of Cognitive Science*. Cambridge MA: MIT Press.

Fodor, J. (1987). *Psychosemantics: The Problem of Meaning in the Philosophy of Mind*. Cambridge MA: MIT Press.

Francescotti, R. (1995). Higher-Order Thoughts and Conscious Experience. *Philosophical Psychology 8*: 239–54.

Gallagher, S. and F.J. Varela (2003). Redrawing the Map and Resetting the Time: Phenomenology and the Cognitive Sciences. *Canadian Journal of Philosophy Supplementary Volume 29:* 93–132.

Ganeri, J. (1999). Self-Intimation, Memory and Personal Identity. *Journal of Indian Philosophy 27:* 469–83.

Gennaro, R. (1996). *Consciousness and Self-Consciousness: A Defense of the Higher-order Thought Theory of Consciousness*. Amsterdam: John Benjamins Publishing Company.

Gennaro, R. (1999). Leibniz on Consciousness and Self-Consciousness. In R. Gennaro and C. Huenemann (Eds.), *New Essays on the Rationalists*, pp. 353–71. New York: Oxford University Press.

Gennaro, R. (2002). Jean-Paul Sartre and the *HOT* Theory of Consciousness. *Canadian Journal of Philosophy 32*: 293–330.

Gennaro, R. (2003). Papineau on the Actualist HOT Theory of Consciousness. *Australasian Journal of Philosophy 81*: 581–6.

Gennaro, R (Ed.) (2004a). *Higher-Order Theories of Consciousness: An Anthology*. Amsterdam: John Benjamins Publishing Company.

Gennaro, R. (2004b). Higher-Order Thoughts, Animal Consciousness, and Misrepresentation: A Reply to Carruthers and Levine. In Gennaro (2004a), pp. 45–66.

Gennaro, R. (2005). The HOT Theory of Consciousness: Between a Rock and a Hard Place. *Journal of Consciousness Studies 12:* 3–21.

Gennaro, R. (2006). Between Pure Self-Referentialism and the Extrinsic HOT Theory. In Kriegel and Williford (2006), pp. 221–48.

Goldman, A. (1993). Consciousness, Folk Psychology, and Cognitive Science. *Consciousness and Cognition 2:* 364–82.

Gurwitsch, A. (1950). *Marginal Consciousness*. Athens OH: Ohio University Press, 1985.

Güzeldere, G. (1995). Is Consciousness the Perception of What Passes in One's Own Mind? In Block, Flanagan, and Güzeldere (1997), pp. 789–806.

Hacker, P.M.S. (1987). *Appearance and Reality: A Philosophical Investigation Into Perception and Perceptual Qualities*. Oxford: Blackwell.

Hacker, P.M.S. (1998). Davidson on the Ontology and Logical Form of Belief. *Philosophy 73:* 81–96.

Hacker, P.M.S. (2001). Eliminative Materialism. In S. Schroeder(Ed.) *Wittgenstein and Contemporary Philosophy of Mind*, pp. 60–84. London: Palgrave.

Hacker, P.M.S. (2002). Is There Anything it is Like to be a Bat? *Philosophy 77:* 157–74.

Hanfling, O. (2000). *Philosophy and Ordinary Language: The Bent and Genius of Our Tongue*. London: Routledge.

Hanfling, O. (2001). Consciousness: 'The Last Mystery.' In S. Schroeder(Ed.) *Wittgenstein and Contemporary Philosophy of Mind*, pp. 36–59. London: Palgrave.

Hardin, C.L. (1988). *Color for Philosophers: Unweaving the Rainbow*. Indianapolis: Hackett.

Harman, G. (1990). The Intrinsic Quality of Experience. In J. Tomberlin (Ed.), *Philosophical Perspectives 4*: 31–52. Atascadero: Ridgeview.

Harman, G. (1996). Explaining Objective Color in Terms of Subjective Experience. In Enrique Villanueva (Ed.), *Philosophical Issues 7*: 1–17. Atascadero: Ridgeview.

Harrison, P. (1991). Do Animals Feel Pain? *Philosophy 66*: 25–40.

Hellie, B. (2007a). Higher-Order Intentionality and Higher-Order Acquaintance. *Philosophical Studies 134*: 289–324.

Hellie, B. (2007b). 'There is something it is like' and the Structure of Consciousness. *Philosophical Review 116*: 441–63.

Henrich, D. (1982). Fichte's Original Insight. In *Contemporary German Philosophy 1*: 15–53.

Hill, C. (2004). "Ouch!" An Essay on Pain. In Gennaro (2004a), pp. 337–62.

Hofstadter, D. and D. Dennett, eds. (1981). *The Mind's I: Fantasies and Reflections on Self and Soul*. Toronto: Bantam Books.

Honderich, T. (1988). *Mind and Brain: A Theory of Determinism, Volume 1*. Oxford: Clarendon.

Hossack, K. (2002). Self-Knowledge and Consciousness. *Proceedings of the Aristotelian Society 102*: 163–81.

Humberstone, I.L. (1996). Intrinsic/Extrinsic. *Synthese 108*: 205–67.

Hume, D. (1978). *A Treatise of Human Nature*. 2nd ed. Ed. L.A. Selby-Bigge. Oxford: Clarendon.

Husserl, E. (1960). *Cartesian Meditations: An Introduction to Phenomenology*. [Trans. Dorion Cairns]. Dordrecht: Martinus Nijhoff.

Jackson, F. (2004). Forward: Looking Back on the Knowledge Argument. In P. Ludlow, D. Stoljar, and Y. Nagasawa (Eds.), *There's Something about Mary*, pp. xv–ix. Cambridge MA: MIT Press.

Jacob, P. (1996). State Consciousness Revisited. *Acta Analytica 16/17*: 29–54.

James, W. (1890). *The Principles of Psychology*. New York: Holt.

Janzen, G. (2005). Self-Consciousness and Phenomenal Character. *Dialogue 44*: 707–33.

Janzen, G. (2006). The Representational Theory of Phenomenal Character: A Phenomenological Critique. *Phenomenology and the Cognitive Sciences 5*: 321–39.

Janzen, G. (2006). Phenomenal Character as Implicit Self-Awareness. *Journal of Consciousness Studies 13*: 44–73.

Jaynes, J. (1976). *The Origin of Consciousness in the Breakdown of the Bicameral Mind*. Boston: Houghton Mifflin Company.

Johnston, M. (1992). How to Speak of the Colors. *Philosophical Studies 68:* 221–63.

Kant, I. (1965). *Critique of Pure Reason*. [Trans. N. Smith]. New York: St. Martin's Press.

Kapitan, T. (1999). The Ubiquity of Self-Awareness. *Grazer Philosophische Studien 57*: 17–44.

Kenny, A. (1984). The Homunculus Fallacy. *The Legacy of Wittgenstein*, pp. 125–36. Oxford: Blackwell.

Kim, J. (1998). *Philosophy of Mind*. Boulder: Westview Press.

Kind, A. (2003). What's so Transparent About Transparency? *Philosophical Studies 115*: 225–44.

Kobes, B.W. (1995). Telic Higher-Order Thoughts and Moore's Paradox. In J. Tomberlin (Ed.). *Philosophical Perspectives 9*: 291–312. Atascadero: Ridgeview.

Kriegel, U. (2003a). Consciousness as Sensory Quality and as Implicit Self-Awareness. *Phenomenology and the Cognitive Sciences 2*: 1–26.

Kriegel, U. (2003b). Consciousness as Intransitive Self-Consciousness: Two Views and an Argument. *Canadian Journal of Philosophy 33*: 103–32.

Kriegel, U. (2003c). Intrinsic Theory and the Content of Inner Awareness. *Journal of Mind and Behavior 24*: 169–96.

Kriegel, U. (2004). Consciousness and Self-Consciousness. *The Monist 87*: 182–205.

Kriegel, U. (2005). Naturalizing Subjective Character. *Philosophy and Phenomenological Research 71*: 23–57.

Kriegel, U. and K. Williford (Eds.) (2006). *Self-Representational Approaches to Consciousness*. Cambridge MA: MIT Press.

Langsam, H. (2000). Experiences, Thoughts, and Qualia. *Philosophical Studies 99*: 269–95.

Levi, D. (1997). Representation: The Eleventh Problem of Consciousness. *Inquiry 40*: 457–73.

Levine, J. (2001). *Purple Haze: The Puzzle of Consciousness*. Oxford: Oxford University Press.

Lewis, D. (1999). Reduction of Mind. In *Papers in Metaphysics and Epistemology*, pp. 291–324. Cambridge: Cambridge University Press.

Loar, B. (2003). Transparent Experience and the Availability of Qualia. In Q. Smith and A. Jokic (Eds.), *Consciousness: New Philosophical Perspectives*, pp. 77–96. Oxford: Clarendon.

Locke, J. (1975). *An Essay Concerning Human Understanding*. Ed. P. Nidditch. Oxford: Clarendon.

Lormand, E. (2004). The Explanatory Stopgap. *Philosophical Review 113:* 303–57.

Lurz, R. (2003). Advancing the Debate Between HOT and FO Accounts of Consciousness. *Journal of Philosophical Research 28*: 23–44.

Lurz, R. (2004). Either FOR or HOR: A False Dichotomy. In Gennaro (2004a), pp. 227–54.

Lycan, W.G. (1990). What is the 'Subjectivity' of the Mental? In J. Tomberlin (Ed.), *Philosophical Perspectives 4*: 109–30. Atascadero: Ridgeview.

Lycan, W.G. (1996). *Consciousness and Experience*. Cambridge MA: MIT Press.

Lycan, W.G. (2001). A Simple Argument for a Higher-Order Representation Theory of Consciousness. *Analysis 61*: 3–4.

Lycan, W.G. (2004). The Superiority of HOP to HOT. In Gennaro (2004a), pp. 93–113.

Lycan, W.G. and Z. Ryder (2003). The Loneliness of the Long-Distance Truck Driver. *Analysis 63:* 132–6.

Malcolm, N. (1977). *Thought and Knowledge*. Ithaca: Cornell University Press.

Malcolm, N. (1984). The Subjective Character of Experience. In Armstrong and Malcolm (1984), pp. 45–66.

Malcolm, N. (1991). I Believe that *p*. In E. Lepore and R. Van Gulick (Eds.), *John Searle and His Critics*, pp. 159–68. Oxford: Basil Blackwell.

Manson, N.C. (2002). What Does Language Tell Us About Consciousness? First-Person Mental Discourse and Higher-Order Thought Theories of Consciousness. *Philosophical Psychology 15*: 221–38.

Martin, C.B. (1997). On the Need for Properties: The Road to Pythagoreanism and Back. *Synthese 112*: 193–231.

Martin, C.B. and K. Pfeifer (1986). Intentionality and the Non-Psychological. *Philosophy and Phenomenological Research 46*: 531–54.

Maund, B. (1995). *Colours: Their Nature and Representation*. Cambridge: Cambridge University Press.

McBride, R. (1999). Consciousness and the State/Transitive/Creature Distinction. *Philosophical Psychology 12:* 181–96.

McCulloch, G. (1999). Bipartism and the Phenomenology of Content. *The Philosophical Quarterly 49*: 18–32

McDowell, J. (1994). *Mind and World*. Cambridge, MA: Harvard University Press.

McGinn, C. (1982). *The Character of Mind*. Oxford: Oxford University Press.

McGinn, C. (1983). *The Subjective View*. Oxford: Clarendon.

McGinn, C. (1989). *Mental Content*. Oxford: Blackwell.

McGinn, C. (1991a). Can We Solve the Mind-Body Problem? In *The Problem of Consciousness: Essays Towards a Resolution*, pp. 1–22. Oxford: Blackwell.

McGinn, C. (1991b). Consciousness and Content. In *The Problem of Consciousness: Essays Towards a Resolution*, pp. 23–43. Oxford: Blackwell.

McGinn, C. (1991c). Could a Machine Be Conscious? In *The Problem of Consciousness: Essays Towards a Resolution*, pp. 202–13. Oxford: Blackwell.

McGinn, C. (1995). Consciousness Evaded: Comments on Dennett. In J. Tomberlin (Ed.) *PhilosophicalPerspectives 9*: 241–9. Atascadero: Ridgeview.

McGinn, C. (1996). Another Look at Color. *Journal of Philosophy 93:* 537–53.

McGinn, C. (2003). What Constitutes the Mind-Body Problem? In E. Sosa and E. Villanueva (Eds.), *Philosophical Issues 13:* 148–62. Oxford: Blackwell.

Merleau-Ponty, M. (1962). *Phenomenology of Perception*. [Trans. Colin Smith]. London: Routledge.

Milner, A.D. and M. Goodale (1995). *The Visual Brain in Action*. Oxford: Oxford University Press.

Moore, G.E. (1922). The Refutation of Idealism. In *Philosophical Studies*. London: Routledge.

Moran, R. (2001). *Authority and Estrangement: An Essay on Self-Knowledge*. Princeton: Princeton University Press.

Nagel, T. (1965). Physicalism. *Philosophical Review 74:* 339–56.

Nagel, T. (1974). What is it Like to Be a Bat? *Philosophical Review 83*: 435–50.

Natsoulas, T. (1992). Toward an Improved Understanding of Sigmund Freud's Conception of Consciousness. *Journal of Mind and Behaviour 13*: 171–92.

Natsoulas, T. (1993). What is Wrong with Appendage Theory of Consciousness? *Philosophical Psychology 6*: 137–54.

Natsoulas, T. (2001). On the Intrinsic Nature of States of Consciousness. *Journal of Mind and Behaviour 22*: 219–48.

Natsoulas, T. (2004). The Case for Intrinsic Theory: X. A Phenomenologist's Account of Inner Awareness. *Journal of Mind and Behaviour 25*: 97–122.

Neander, K. (1998). The Division of Phenomenal Labor: A Problem for Representational Theories of Consciousness. In J. Tomberlin (Ed.), *Philosophical Perspectives 12*: 411–34. Atascadero: Ridgeview.

Nelkin, N. (1986). Pains and Pain Sensations. *Journal of Philosophy 83*: 129–48.

Nelkin, N. (1987). What is it Like to be a Person? *Mind & Language 2*: 220–41.

Nelkin, N. (1989a). Unconscious Sensations. *Philosophical Psychology 2*: 129–41.

Nelkin, N. (1989b). Propositional Attitudes and Consciousness. *Philosophy and Phenomenological Research 49*: 413–30.

Nelkin, N. (1993). What is Consciousness? *Philosophy of Science 60*: 419–34.

Papineau, D. (2002). *Thinking About Consciousness*. Oxford: Clarendon.

Peacocke, C. (1983). *Sense and Content*. Oxford: Oxford University Press.

Peacocke, C. (1992). *A Study of Concepts*. Cambridge: MIT Press.

Penrose, R. (1989). *The Emperor's New Mind*. Oxford: Oxford University Press.

Perry, J. (1979). The Problem of the Essential Indexical. *Noûs 13*: 3–21.

Pitt, D. (2004). The Phenomenology of Cognition, Or *What is it Like to Think That P*? *Philosophy and Phenomenological Research 69*: 1–36.

Place, U.T. (1956). Is Consciousness a Brain Process? *British Journal of Psychology 47*: 44–50.

Quine, W. (1951). Two Dogmas of Empiricism. *Philosophical Review* **60**: 20–43.

Quine, W. (1960). *Word and Object*. Cambridge, MA: MIT Press.

Rey, G. (1988). A Question About Consciousness. In H.R. Otto and J.A. Tuedio (Eds), *Perspectives on Mind*, pp. 5–24. Dordrecht: D. Reidel Publishing Company.

Rey, G. (1998). A Narrow Representationalist Account of Qualitative Experience. In J. Tomberlin(Ed.), *Philosophical Perspectives 12*: 435–57. Atascadero: Ridgeview.

Robinson, W. (2004). A Few Thoughts Too Many? In Gennaro (2004a), pp. 295–313.

Robinson, H. (1994). *Perception*. London: Routledge.

Rorty, R. (1979). *Philosophy and the Mirror of Nature*. Princeton: Princeton University Press.

Rosenthal, D. (1986). Two Concepts of Consciousness. *Philosophical Studies 49:* 329–59.

Rosenthal, D. (1990). A Theory of Consciousness. In Block, Flanagan, and Güzeldere (1997), pp. 729–53.

Rosenthal, D. (1993a). State Consciousness and Transitive Consciousness. *Consciousness and Cognition 2*: 355–63.

Rosenthal, D. (1993b). Higher-Order Thoughts and the Appendage Theory of Consciousness. *Philosophical Psychology 6*: 155–67.

Rosenthal, D. (1993c). Multiple Drafts and Higher-Order Thoughts. *Philosophy and Phenomenological Research 53*: 911–18.

Rosenthal, D. (1993d). Thinking that one Thinks. In M. Davies and G.W. Humphreys (Eds.), *Consciousness: Psychologicaland Philosophical Essays*, pp. 197–223. Oxford: Blackwell.

Rosenthal, D. (2000a). Consciousness, Content, and Metacognitive Judgments. *Consciousness and Cognition 9:* 203–14.

Rosenthal, D. (2000b). Consciousness and Metacognition. In D. Sperber (Ed.), *Metarepresentation: A Multidisciplinary Perspective*, pp. 265–95. New York: Oxford University Press.

Rosenthal, D. (2002a). How Many Kinds of Consciousness? *Consciousness and Cognition 11*: 653–65.

Rosenthal, D. (2002b). Consciousness and the Mind. *Iyyun 51*: 227–51.
Rosenthal, D. (2002c). Explaining Consciousness. In D. Chalmers (Ed.), *Philosophy of Mind: Classicaland Contemporary Readings*, pp. 406–21. New York: Oxford University Press.
Rosenthal, D. (2002d). Consciousness and Higher-Order Thought. In L. Nadel(Ed.), *Encyclopedia of Cognitive Science*, pp. 717–26. London: Macmillan.
Rosenthal, D. (2004a). Varieties of Higher-Order Theory. In Gennaro (2004a), pp. 17–44.
Rosenthal, D. (2004b). Being Conscious of Ourselves. *The Monist 87:* 159–81.
Rosenthal, D. (2005). *Consciousness and Mind*. Oxford: Clarendon.
Rowlands, M. (2001). *The Nature of Consciousness*. Cambridge: Cambridge University Press.
Ryle, G. (1949). *The Concept of Mind*. Chicago: University of Chicago Press.
Sartre, J.P. (1956). *Being and Nothingness*. Trans. H. Barnes. New York: Philosophical Library.
Sartre, J.P. (1957). *The Transcendence of the Ego: An Existentialist Theory of Consciousness*. [Trans. F. Williams and R. Kirkpatrick]. New York: Farrar, Straus and Company.
Seager, W. (1999). *Theories of Consciousness: An Introduction and Assessment*. London: Routledge.
Seager, W. (2004). A Cold Look at HOT Theory. In Gennaro (2004a), pp. 255–75.
Searle, J. (1979). What is an Intentional State? *Mind 88:* 74–92.
Searle, J. (1983). *Intentionality: An Essay in the Philosophy of Mind*. Cambridge: Cambridge University Press.
Searle, J. (1984). *Minds, Brains and Science*. Cambridge MA: Harvard University Press.
Searle, J. (1992). *The Rediscovery of the Mind*. Cambridge MA: MIT Press.
Searle, J. (1997). *The Mystery of Consciousness*. New York: New York Review of Books.
Searle, J. (1999). *Mind, Language and Society: Philosophy in the Real World*. London: Weidenfeld & Nicolson.
Sellars, W. (1997). *Empiricism and the Philosophy of Mind*. Cambridge: Harvard University Press.
Schröder, J. (2001). Higher-Order Thought and Naturalist Accounts of Consciousness. *Journal of Consciousness Studies 8*: 27–46.
Shoemaker, S. (1994). Phenomenal Character. *Noûs 28*: 21–38.
Shoemaker, S. (1996a). Qualia and Consciousness. In *The First-Person Perspective and other Essays*, pp. 121–40. Cambridge: Cambridge University Press.
Shoemaker, S. (1996b). Self-Knowledge and 'Inner Sense', Lecture II: The Broad Perceptual Model. In *The First-Person Perspective and Other Essays*, pp. 224–45. Cambridge: Cambridge University Press.
Shoemaker, S. (1996c). Self-Knowledge and 'Inner Sense', Lecture III: The Phenomenal Character of Experience. In *The First-Person Perspective and Other Essays*, pp. 246–68. Cambridge: Cambridge University Press.
Siewert, C. (1998). *The Significance of Consciousness*. Princeton: Princeton University Press.
Sprigge, T.L. S. (1982). The Importance of Subjectivity. *Inquiry 25:* 143–63.
Stalnaker, R. (1999). *Context and Content*. Oxford: Oxford University Press.
Strawson, G. (1994). *Mental Reality*. Cambridge MA: MIT Press.
Stroud, B. (2000). *The Quest for Reality: Subjectivism and the Metaphysics of Colour*. New York: Oxford University Press.
Stubenberg, L. (1998). *Consciousness and Qualia*. Amsterdam: John Benjamins Publishing Company.
Thomas, A. (2003). An Adverbial Theory of Consciousness. *Phenomenology and and the Cognitive Sciences 2*: 161–85.

Thomasson, A. (2000). After Brentano: A One-Level Theory of Consciousness. *European Journal of Philosophy 8:* 190–209.

Thompson, B. (2007). Shoemaker on Phenomenal Content. *Philosophical Studies 135:* 307–34.

Tye, M. (1995). *Ten Problems of Consciousness: A Representational Theory of the Phenomenal Mind*. Cambridge MA: MIT Press.

Tye, M. 1996. The Function of Consciousness. *Noûs 30:* 287–305.

Tye, M. (2000). *Consciousness, Color, and Content*. Cambridge MA: MIT Press.

Van Gulick, R. (1995). How Should We Understand the Relation Between Intentionality and Phenomenal Consciousness? In J. Tomberlin (Ed.), *Philosophical Perspectives 9*: 271–89. Atascadero: Ridgeview.

Van Gulick, R. (2004). Higher-Order Global States (HOGS): An Alternative Higher-Order Model of Consciousness. In Gennaro (2004a), pp. 67–92.

Van Gulick, R. (2005). Inward and Upward: Reflection, Introspection, and Self-Awareness. *Philosophical Topics 28*: 275–305.

Van Gulick, R. (2006). Mirror Mirror–Is That All? In Kriegel and Williford (2006), pp. 11–39.

White, A.R. (1964). *Attention*. Oxford: Blackwell.

White, S. (1987). What is it Like to Be an Homunculus? *Pacific Philosophical Quarterly 68:* 148–174.

Wilkes, K. (1984). Is Consciousness Important? *British Journal for the Philosophy of Science 35*: 223–43.

Williamson, T. (2000). *Knowledge and Its Limits*. Oxford: Oxford University Press.

Wittgenstein, L. (1953). *Philosophical Investigations*. 2nd ed. Eds. G.E. M. Anscombe and R. Rhees. [Trans. G.E. M. Anscombe]. Oxford: Blackwell.

Wittgenstein, L. (1967). *Zettel*. 2nd. ed. Eds. G.E. M. Anscombe and G.H. von Wright. Trans. G.E. M. Anscombe. Oxford: Blackwell.

Wittgenstein, L. (1980). *Remarks on the Philosophy of Psychology, Volume II*. Eds. G.H. von Wright and H. Nyman. Trans. C.G. Luckhardt and M.A. E. Aue. Oxford: Blackwell.

Woodruff Smith, D. (1986). The Structure of (Self-)Consciousness. *Topoi 5:* 149–56.

Woodruff Smith, D. (1988). Rey Cogitans: The Unquestionability of Consciousness. In H.R. Otto and J.A. Tuedio (Eds.), *Perspectives on Mind*, pp. 25–34. Dordrecht: D. Reidel Publishing Company.

Woodruff Smith, D. (1989). *The Circle of Acquaintance: Perception, Consciousness, and Empathy*. Dordrecht: Kluwer.

Woodruff Smith, D. (1993). The Cogito circa AD 2000. *Inquiry 36*: 225–54.

Woodruff Smith, D. (2004a). Three Facets of Consciousness. In *Mind World: Essays in Phenomenology and Ontology*, pp. 10–41. Cambridge: Cambridge University Press.

Woodruff Smith, D. (2004b). Return to Consciousness. In *Mind World: Essays in Phenomenology and Ontology*, pp. 76–121. Cambridge: Cambridge University Press.

Worley, S. (2003). Conceivability, Possibility and Physicalism. *Analysis 63:* 15–23.

Zahavi, D. (1999). *Self-Awareness and Alterity: A Phenomenological Investigation*. Evanston: Northwestern University Press.

Zahavi, D. (2002a). First-Person Thoughts and Embodied Self-Awareness: Some Reflections on the Relation Between Recent Analytical Philosophy and Phenomenology. *Phenomenology and the Cognitive Sciences 1*: 7–26.

Zahavi, D. (2002b). The Three Concepts of Consciousness in *Logische Untersuchungen*. *Husserl Studies 18:* 51–64.

Zahavi, D. (2003). Inner Time-Consciousness and Pre-Reflective Self-Awareness. In D. Welton(Ed.), *The New Husserl: A Critical Reader*, pp. 157–80. Bloomington: Indiana University Press.

Zahavi, D. (2004). Back to Brentano? *Journal of Consciousness Studies 11*: 66–87.

Zahavi, D. (2005). *Subjectivity and Selfhood: Investigating the First-Person Perspective*. Cambridge MA: MIT Press.

Zahavi, D. (2007). Killing the Straw Man: Dennett and Phenomenology. *Phenomenology and the Cognitive Sciences 6*: 21–43.

Zahavi, D. and J. Parnas (1998). Phenomenal Consciousness and Self-Awareness: A Phenomenological Critique of Representational Theory. *Journal of Consciousness Studies 5*: 687–705.

Index of subjects

A
access consciousness 35–7, 45
acquaintance 116, 132, 172
animal consciousness 13, 14, 23–4, 127–30
appearance/reality distinction 7, 9

B
behaviourism 20, 171
beliefs 8, 13–14, 16, 22, 37, 39, 64, 70, 78–82, 88, 134
  unconscious 21–2
blindsight 35, 39

C
concepts 13, 20, 31, 87, 124, 127
conceptual truth 24, 44–5, 56
confabulated conscious states 96–9, 115
consciousness, evolutionary value of 168–9
content 15, 125, 128, 134, 135, 140, 142
  see also de se content
  sub-personal level content
  personal level content
creature consciousness 13, 17–18, 46, 57

D
de se content 106, 118–21, 124
desires 22
  unconscious 22
dispositions 66, 68
dreaming consciousness 14, 15, 18, 48

E
Empiricism 38, 151, 157, 159
eliminativism 8
emotions 65–6
experience error 159

F
fact awareness 13, 71
folk psychology 5–6
focal awareness 21, 52, 116

G
generation problem 133
generality problem 93

H
higher-order (HO) theory of consciousness 3, 50–1, 86f.
  dispositionalist version of 87, 88, 95
  actualist version of 88

I
imagination 105, 156, 167
implicit self-awareness 2, 4, 102, 106–7, 112, 115–18, 121–2, 124–9, 131, 135, 145–7, 154–6, 160–1, 165–7, 169, 172
infant consciousness 14, 19, 23–4, 55, 69, 83, 95, 124, 127–31
intensionality 62–4
intentionality 3, 15–16, 59–60, 62, 64–9, 103, 135, 138
intransitive creature consciousness 3, 13–14
intransitive state consciousness 20, 34, 89, 96
introspection 22, 34, 146, 161–2, 167

M
mereological principle 18, 20
moods 16, 59, 66–8, 144
Müller-Lyer illusion 73

N
non-objectual acquaintance, *see* acquaintance

O
object awareness 116–17

P
pain 7, 9, 21, 23, 25, 46, 52–6, 60–5, 68, 90, 95–97, 99, 123, 135, 146, 157–8, 162, 167–71
  congenital insensitivity to 168
  evolutionary value of 168
  location of 62–4
  peripheral awareness of 21
  unconscious 21
perception 6–7, 23, 34, 41–5, 48, 70, 79, 84, 87, 107, 109–11, 113, 117, 120–3, 126–7, 140, 143, 151–3, 158–9, 161–2, 167–9
peripheral awareness 21, 47–8, 51–2, 116
personal-level content 5–8, 30, 73
phenomenal character 1–5, 9, 12, 24–8, 34, 36–7, 39–41, 45–7, 49, 51–3, 56–7, 59, 77, 135, 137–8, 140–2, 144, 147, 153–7, 159–67, 169–70, 172
  reductive explanation of 155–6
phenomenal consciousness 2–3, 35–7, 45
  definition of 23
phenomenology 3–9, 72, 99–102, 113–18, 126, 128, 138, 141, 148–9, 152, 157–9, 163, 165, 172
physicalism 2, 39
privacy 27, 169–71
privileged access 170–1
propositional attitudes 37, 38, 54, 114

property assumption 141, 144, 147–8, 150, 153–4

Q
qualia 25–9, 114, 161–3, 165–6, 169–70
qualitative feels 25, 29, 54
see also qualia

R
regress argument 104, 110
representationalism 65, 135, 137, 140–2, 147, 151, 153–4

S
self, the 129–31
self-awareness 1–4, 17, 57, 59, 69, 77–9, 83, 85–6, 98, 100–3, 105–7, 111–18, 128–30, 155, 162, 172
and self-consciousness 1, 76
self-knowledge 134–5
semantic ascent 17
sensations 8, 16, 20, 23, 25, 46, 53–5, 60, 62, 65, 67, 71, 75, 121, 150, 156–8, 166–7, 171–2
spontaneous reportability 73–5, 79–83
state consciousness 2–4, 17–18, 28, 33–6, 41–6, 53, 57, 59, 69, 77, 79, 84, 92, 94, 103–4, 106, 115, 119, 132–3, 145, 160
subject-object relation 4, 77–8, 102, 115–17
see also object awareness
sub-personal level content 6–8, 32, 44, 73, 133, 152
symmetry argument 77–8, 83

T
thing awareness 13, 71
transparency assumption 141–3, 147, 153–4
transitive creature consciousness 3, 13–14, 20, 34, 36, 89, 96, 120
transitivity principle 69, 74–5, 79, 83, 106, 112–13, 173

U
unconscious awareness 1, 15, 71, 77, 102, 113, 146
unconscious mental states 16, 18, 43, 71, 93, 98, 104, 106, 111–12, 114
see also pains, unconscious
beliefs, unconscious
desires, unconscious

W
"what-it-is-like" locution 3, 24–5, 27–31, 33
what-it-is-likeness 17, 28–30, 35, 46, 54–6, 59, 155, 157, 160–3, 167, 170, 172

Z
zombies 39, 168

Index of names

A
Antony, M. 11
Aquila, R. 85, 113
Armstrong, D.M. 13, 15, 36, 46, 50–1, 86–7, 96, 113, 125, 162

B
Baier, A. 68
Baker , L.R. 85, 130
Balog, K. 39
Bennett, J. 1, 18, 20, 22, 37, 60, 66, 69, 80–1, 117, 123–4, 133, 138, 158, 162, 168–9, 171
Bennett, M.R. 1, 18, 20, 22, 37, 60, 66, 69, 80–1, 117, 123–4, 133, 138, 158, 162, 168–9, 171
Bermúdez, J. 130–1
Block, N. 25–6, 35–40, 42, 44–5, 62–4, 85, 137
Boghossian, P. 8, 157
Bower, T.G.R. 131
Brentano, F. 1–5, 13, 16, 59–60, 71, 74, 86, 103–15, 117–18, 125–8, 132, 138, 154
Broad, C.D. 149
Brueckner, A. 39
Burge, T. 36
Byrne, A. 16, 36, 52, 73, 85, 88, 93, 96, 137, 141, 164

C
Campbell, S. 149
Carman, T. 4
Carruthers, P. 13, 17–18, 21, 24, 26, 35, 39, 45–6, 55, 62, 87–8, 95, 99–100, 113, 125, 127, 141, 157, 164–5
Castañeda, H.-N. 118
Caston, V. 96, 104, 109–10, 113, 132, 135
Casullo, A. 143
Chalmers, D. 2, 11, 19, 25–6, 36, 39, 85, 116, 133, 157, 163, 168
Chisholm, R. 75, 118
Churchland, P.M. 19, 133
Clark, A. 44
Collins, A. 79–81, 155
Cowley, F. 5, 18
Crane, T. 2, 16, 60, 137
Crick, F. 133

D
Dainton, B. 4
Damasio, A. 65–6
Dennett, D. 6–8, 19, 24–7, 40–1, 52, 55, 83, 86, 93, 104, 168
Descartes 19, 23–4, 39, 113, 157
Douglas, G. 60, 117, 169
Dretske, F. 13–14, 25–6, 60, 69–71, 74, 77, 85–6, 93, 95, 135, 137–40, 142, 147–8, 153, 168
Droege, P. 85
Drummond, J.J. 115
Dwyer, P. 4, 6–7, 48–9, 70, 151–2, 167

E
Edelman, G. 2, 6

F
Farrell, B. 2
Fichte, J. 113
Field, J. 130
Flanagan, O. 2, 8, 40, 54, 155
Fodor, J. 64, 172
Francescotti, R. 85

G
Gallagher, S. 5
Ganeri, J. 75
Gennaro, R. 1, 3, 25–6, 43, 53, 56, 85–8, 91, 94–95, 97–101, 104, 108, 113, 118–19, 125, 132, 162
Goldman, A. 8, 85, 93
Goodale, M. 18
Gurwitsch, A. 2, 73, 106, 113
Güzeldere, G. 17, 85, 104

H
Hacker , P.M.S. 1, 8, 18, 20, 22, 24–6, 29–32, 37, 60, 66, 77, 79–81, 117, 123–4, 133, 138, 157–8, 162, 168–9, 171
Hanfling, O. 24, 45
Hardin, C.L. 157
Harman, G. 137, 141, 143, 153, 158
Harrison, P. 24
Heidegger 113, 116
Hellie, B. 30, 85
Henrich, D. 176
Hill, C. 9
Hofstadter, D. 24
Honderich, T. 15
Hossack, K. 9, 85, 134–5
Humberstone, I.L. 104
Hume, D. 22–3
Husserl, E. 4–5, 77, 113, 115–16

J
Jackson, F. 137
Jacob, P. 85
James, W. 113
Janzen, G. 177
Jaynes, J. 41
Johnston, M. 157

K
Kant, I. 19, 33, 71, 87, 113
Kapitan, T. 119
Kenny, A. 18

Kim, J. 64
Kind, A. 142–3
Kobes, B.W. 92
Koch, C. 133
Kriegel, U. 2, 11, 43, 96, 100–1, 111–12, 116, 119–20, 122, 156

L
Langsam, H. 163
Leibniz 71, 113
Levi, D. 24
Levine, J. 39, 85, 96, 133, 156
Lewis, D. 118
Loar, B. 159
Locke, J. 87, 113, 161
Lormand, E. 30, 32
Lurz, R. 17, 85, 95
Lycan, W.G. 3, 11, 24–5, 27–9, 52, 64, 69, 86–7, 96, 101, 113, 125, 137, 139, 143, 148, 157, 162, 170

M
Malcolm, N. 13–15, 24, 80, 171
Manson, N.C. 85
Martin, C.B. 61, 67–8
Maund, B. 157
McBride, R. 11
McCulloch, G. 163
McDowell, J. 124
McGinn, C. 2, 7–8, 19, 41, 60, 116, 118, 155, 157
Mellor, D.H. 2
Merleau-Ponty, M. 4–6, 15, 70, 107, 116, 152–4, 157, 159
Milner, A.D. and 18
Moore, G.E. 141
Moran, R. 162

N
Nagel, T. 2, 4–5, 7, 23–4, 28–33, 41, 46
Natsoulas, T. 85, 109, 111, 113
Neander, K. 85, 96
Nelkin, N. 11, 35, 45–6, 52–5, 86, 157, 161

P
Papineau, D. 85
Parnas, J. 2, 5, 85
Peacocke, C. 73, 159
Penrose, R. 133
Perry, J. 118
Pfeifer, K. 61, 67–8
Pitt, D. 163
Place, U.T. 159

Q
Quine, W. 17, 24

R
Rey, G. 8, 64, 137–8
Robinson, W. 85, 143
Robinson, H. 143
Rorty, R. 14
Rosenthal, D. 3, 12–13, 17, 21, 36, 59, 69, 73–4, 83, 86–98, 100–1, 104–5, 113, 118, 125, 131–2, 141, 162, 165
Rowlands, M. 6–7, 85, 93, 96, 141
Ryle, G. 22, 45, 117, 133, 152

S
Sartre, J.P. 1–2, 4, 73–5, 77, 80, 103–4, 113, 115–18, 154, 163
Schröder, J. 85, 132
Seager, W. 85, 95–96, 133, 138
Searle, J. 7–8, 11, 19, 22, 46, 48–9, 51, 60, 64, 68–9, 106, 162, 170
Sellars, W. 13–14
Shoemaker, S. 9, 22, 25–6, 137, 141, 162
Siewert, C. 4, 68, 85, 104
Sprigge, T.L.S. 2
Stalnaker, R. 73
Strawson, G. 163
Stroud, B. 157
Stubenberg, L. 4, 7, 85, 93

T
Thomas, A. 115
Thomasson, A. 4, 85, 93, 111, 115, 117
Thompson, B. 156
Tye, M. 16, 60–7, 77, 137–45, 147–53, 156–7, 164, 169

V
Van Gulick, R. 17, 93, 113
Varela, F.J. 5
Velleman, J.D. 157

W
White, A.R. 1, 18–19, 69–70, 72, 106
White, S. 18, 70, 106
Wilkes, K. 8, 11
Williford, K. 177
Williamson, T. 79
Wittgenstein, L. 20, 45, 80–1, 104, 123, 152–3, 162, 171
Woodruff Smith, D. 4, 8, 15, 69, 73, 85, 93, 100, 106, 111, 113–14, 116, 119–20, 122, 155, 162–3
Worley, S. 39

Z
Zahavi, D. 2, 4–6, 85, 100, 104, 113, 115–16, 120, 123, 154, 163

Advances in Consciousness Research

A complete list of titles in this series can be found on the publishers' website, *www.benjamins.com*

72 **JANZEN, Greg:** The Reflexive Nature of Consciousness. 2008. vii, 186 pp.

71 **KROIS, John Michael, Mats ROSENGREN, Angela STEIDELE and Dirk WESTERKAMP (eds.):** Embodiment in Cognition and Culture. 2007. xxii, 304 pp.

70 **RAKOVER, Sam S.:** To Understand a Cat. Methodology and philosophy. 2007. xviii, 253 pp.

69 **KUCZYNSKI, John-Michael:** Conceptual Atomism and the Computational Theory of Mind. A defense of content-internalism and semantic externalism. 2007. x, 524 pp.

68 **BRÅTEN, Stein (ed.):** On Being Moved. From mirror neurons to empathy. 2007. x, 333 pp.

67 **ALBERTAZZI, Liliana (ed.):** Visual Thought. The depictive space of perception. 2006. xii, 380 pp.

66 **VECCHI, Tomaso and Gabriella BOTTINI (eds.):** Imagery and Spatial Cognition. Methods, models and cognitive assessment. 2006. xiv, 436 pp.

65 **SHAUMYAN, Sebastian:** Signs, Mind, and Reality. A theory of language as the folk model of the world. 2006. xxvii, 315 pp.

64 **HURLBURT, Russell T. and Christopher L. HEAVEY:** Exploring Inner Experience. The descriptive experience sampling method. 2006. xii, 276 pp.

63 **BARTSCH, Renate:** Memory and Understanding. Concept formation in Proust's *A la recherche du temps perdu.* 2005. x, 160 pp.

62 **DE PREESTER, Helena and Veroniek KNOCKAERT (eds.):** Body Image and Body Schema. Interdisciplinary perspectives on the body. 2005. x, 346 pp.

61 **ELLIS, Ralph D.:** Curious Emotions. Roots of consciousness and personality in motivated action. 2005. viii, 240 pp.

60 **DIETRICH, Eric and Valerie Gray HARDCASTLE:** Sisyphus's Boulder. Consciousness and the limits of the knowable. 2005. xii, 136 pp.

59 **ZAHAVI, Dan, Thor GRÜNBAUM and Josef PARNAS (eds.):** The Structure and Development of Self-Consciousness. Interdisciplinary perspectives. 2004. xiv, 162 pp.

58 **GLOBUS, Gordon G., Karl H. PRIBRAM and Giuseppe VITIELLO (eds.):** Brain and Being. At the boundary between science, philosophy, language and arts. 2004. xii, 350 pp.

57 **WILDGEN, Wolfgang:** The Evolution of Human Language. Scenarios, principles, and cultural dynamics. 2004. xii, 240 pp.

56 **GENNARO, Rocco J. (ed.):** Higher-Order Theories of Consciousness. An Anthology. 2004. xii, 371 pp.

55 **PERUZZI, Alberto (ed.):** Mind and Causality. 2004. xiv, 235 pp.

54 **BEAUREGARD, Mario (ed.):** Consciousness, Emotional Self-Regulation and the Brain. 2004. xii, 294 pp.

53 **HATWELL, Yvette, Arlette STRERI and Edouard GENTAZ (eds.):** Touching for Knowing. Cognitive psychology of haptic manual perception. 2003. x, 322 pp.

52 **NORTHOFF, Georg:** Philosophy of the Brain. The brain problem. 2004. x, 433 pp.

51 **DROEGE, Paula:** Caging the Beast. A theory of sensory consciousness. 2003. x, 183 pp.

50 **GLOBUS, Gordon G.:** Quantum Closures and Disclosures. Thinking-together postphenomenology and quantum brain dynamics. 2003. xxii, 200 pp.

49 **OSAKA, Naoyuki (ed.):** Neural Basis of Consciousness. 2003. viii, 227 pp.

48 **JIMÉNEZ, Luis (ed.):** Attention and Implicit Learning. 2003. x, 385 pp.

47 **COOK, Norman D.:** Tone of Voice and Mind. The connections between intonation, emotion, cognition and consciousness. 2002. x, 293 pp.

46 **MATEAS, Michael and Phoebe SENGERS (eds.):** Narrative Intelligence. 2003. viii, 342 pp.

45 **DOKIC, Jérôme and Joëlle PROUST (eds.):** Simulation and Knowledge of Action. 2002. xxii, 271 pp.

44 **MOORE, Simon C. and Mike OAKSFORD (eds.):** Emotional Cognition. From brain to behaviour. 2002. vi, 350 pp.

43 **DEPRAZ, Nathalie, Francisco J. VARELA and Pierre VERMERSCH:** On Becoming Aware. A pragmatics of experiencing. 2003. viii, 283 pp.

42 **STAMENOV, Maxim I. and Vittorio GALLESE (eds.):** Mirror Neurons and the Evolution of Brain and Language. 2002. viii, 392 pp.

41 **ALBERTAZZI, Liliana (ed.):** Unfolding Perceptual Continua. 2002. vi, 296 pp.

40 **MANDLER, George:** Consciousness Recovered. Psychological functions and origins of conscious thought. 2002. xii, 142 pp.

39 **BARTSCH, Renate:** Consciousness Emerging. The dynamics of perception, imagination, action, memory, thought, and language. 2002. x, 258 pp.

38 **SALZARULO, Piero and Gianluca FICCA (eds.):** Awakening and Sleep–Wake Cycle Across Development. 2002. vi, 283 pp.

37 **PYLKKÄNEN, Paavo and Tere VADÉN (eds.):** Dimensions of Conscious Experience. 2001. xiv, 209 pp.

36 **PERRY, Elaine, Heather ASHTON and Allan H. YOUNG (eds.):** Neurochemistry of Consciousness. Neurotransmitters in mind. With a foreword by Susan Greenfield. 2002. xii, 344 pp.

35 **Mc KEVITT, Paul, Seán Ó NUALLÁIN and Conn MULVIHILL (eds.):** Language, Vision and Music. Selected papers from the 8th International Workshop on the Cognitive Science of Natural Language Processing, Galway, 1999. 2002. xii, 433 pp.

34 **FETZER, James H. (ed.):** Consciousness Evolving. 2002. xx, 253 pp.

33 **YASUE, Kunio, Mari JIBU and Tarcisio DELLA SENTA (eds.):** No Matter, Never Mind. Proceedings of Toward a Science of Consciousness: Fundamental approaches, Tokyo 1999. 2002. xvi, 391 pp.

32 **VITIELLO, Giuseppe:** My Double Unveiled. The dissipative quantum model of brain. 2001. xvi, 163 pp.

31 **RAKOVER, Sam S. and Baruch CAHLON:** Face Recognition. Cognitive and computational processes. 2001. x, 306 pp.

30 **BROOK, Andrew and Richard C. DEVIDI (eds.):** Self-Reference and Self-Awareness. 2001. viii, 277 pp.

29 **VAN LOOCKE, Philip (ed.):** The Physical Nature of Consciousness. 2001. viii, 321 pp.

28 **ZACHAR, Peter:** Psychological Concepts and Biological Psychiatry. A philosophical analysis. 2000. xx, 342 pp.

27 **GILLETT, Grant R. and John McMILLAN:** Consciousness and Intentionality. 2001. x, 265 pp.

26 **Ó NUALLÁIN, Seán (ed.):** Spatial Cognition. Foundations and applications. 2000. xvi, 366 pp.

25 **BACHMANN, Talis:** Microgenetic Approach to the Conscious Mind. 2000. xiv, 300 pp.

24 **ROVEE-COLLIER, Carolyn, Harlene HAYNE and Michael COLOMBO:** The Development of Implicit and Explicit Memory. 2000. x, 324 pp.

23 **ZAHAVI, Dan (ed.):** Exploring the Self. Philosophical and psychopathological perspectives on self-experience. 2000. viii, 301 pp.

22 **ROSSETTI, Yves and Antti REVONSUO (eds.):** Beyond Dissociation. Interaction between dissociated implicit and explicit processing. 2000. x, 372 pp.

21 **HUTTO, Daniel D.:** Beyond Physicalism. 2000. xvi, 306 pp.

20 **KUNZENDORF, Robert G. and Benjamin WALLACE (eds.):** Individual Differences in Conscious Experience. 2000. xii, 412 pp.

19 **DAUTENHAHN, Kerstin (ed.):** Human Cognition and Social Agent Technology. 2000. xxiv, 448 pp.

18 **PALMER, Gary B. and Debra J. OCCHI (eds.):** Languages of Sentiment. Cultural constructions of emotional substrates. 1999. vi, 272 pp.

17 **HUTTO, Daniel D.:** The Presence of Mind. 1999. xiv, 252 pp.

16 **ELLIS, Ralph D. and Natika NEWTON (eds.):** The Caldron of Consciousness. Motivation, affect and self-organization — An anthology. 2000. xxii, 276 pp.

15 **CHALLIS, Bradford H. and Boris M. VELICHKOVSKY (eds.):** Stratification in Cognition and Consciousness. 1999. viii, 293 pp.

14 **SHEETS-JOHNSTONE, Maxine:** The Primacy of Movement. 1999. xxxiv, 583 pp.

13 **VELMANS, Max (ed.):** Investigating Phenomenal Consciousness. New methodologies and maps. 2000. xii, 381 pp.

12 **STAMENOV, Maxim I. (ed.):** Language Structure, Discourse and the Access to Consciousness. 1997. xii, 364 pp.

11 **PYLKKÖ, Pauli:** The Aconceptual Mind. Heideggerian themes in holistic naturalism. 1998. xxvi, 297 pp.

10 **NEWTON, Natika:** Foundations of Understanding. 1996. x, 211 pp.

9 **Ó NUALLÁIN, Seán, Paul Mc KEVITT and Eoghan Mac AOGÁIN (eds.):** Two Sciences of Mind. Readings in cognitive science and consciousness. 1997. xii, 490 pp.

8 **GROSSENBACHER, Peter G. (ed.):** Finding Consciousness in the Brain. A neurocognitive approach. 2001. xvi, 326 pp.

7 **MAC CORMAC, Earl and Maxim I. STAMENOV (eds.):** Fractals of Brain, Fractals of Mind. In search of a symmetry bond. 1996. x, 359 pp.

6 **GENNARO, Rocco J.:** Consciousness and Self-Consciousness. A defense of the higher-order thought theory of consciousness. 1996. x, 220 pp.

5 **STUBENBERG, Leopold:** Consciousness and Qualia. 1998. x, 368 pp.

4 **HARDCASTLE, Valerie Gray:** Locating Consciousness. 1995. xviii, 266 pp.

3 **JIBU, Mari and Kunio YASUE:** Quantum Brain Dynamics and Consciousness. An introduction. 1995. xvi, 244 pp.

2 **ELLIS, Ralph D.:** Questioning Consciousness. The interplay of imagery, cognition, and emotion in the human brain. 1995. viii, 262 pp.

1 **GLOBUS, Gordon G.:** The Postmodern Brain. 1995. xii, 188 pp.

LIVERPOOL
JOHN MOORES UNIVERSITY
AVRIL ROBARTS LRC
TEL 0151 231 4022

Combining phenomenological insights from Brentano and Sartre, but also drawing on recent work on consciousness by analytic philosophers, this book defends the view that conscious states are reflexive, and necessarily so, i.e., that they have a built-in, "implicit" awareness of their own occurrence, such that the subject of a conscious state has an immediate, non-objectual acquaintance with it. As part of this investigation, the book also explores the relationship between reflexivity and the phenomenal, or "what-it-is-like," dimension of conscious experience, defending the innovative thesis that phenomenal character is constituted by the implicit self-awareness built into every conscious state. This account stands in marked contrast to most influential extant theories of phenomenal character, including qualia theories, according to which phenomenal character is a matter of having phenomenal sensations, and representationalism, according to which phenomenal character is constituted by representational content.

KS-593-140

ISBN 978 90 272 5208 1

9 789027 252081

John Benjamins Publishing Company