Caging the Beast

A theory of sensory consciousness

Paula Droege



Advances in Consciousness Research

Digitized by the Internet Archive
in 2019 with funding from
Kahle/Austin Foundation

https://archive.org/details/cagingbeasttheor0000droe

Caging the Beast

Advances in Consciousness Research

Advances in Consciousness Research provides a forum for scholars from different scientific disciplines and fields of knowledge who study consciousness in its multifaceted aspects. Thus the Series will include (but not be limited to) the various areas of cognitive science, including cognitive psychology, linguistics, brain science and philosophy. 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 and society.

Series A: Theory and Method. Contributions to the development of theory and method in the study of consciousness.

Editor

Maxim I. Stamenov
Bulgarian Academy of Sciences

Editorial Board

David Chalmers
University of Arizona

Gordon G. Globus
University of California at Irvine

Ray Jackendoff
Brandeis University

Christof Koch
California Institute of Technology

Stephen Kosslyn
Harvard University

Earl Mac Cormac
Duke University

George Mandler
University of California at San Diego

John R. Searle
University of California at Berkeley

Petra Stoerig
Universität Düsseldorf

† Francisco Varela
C.R.E.A., Ecole Polytechnique, Paris

Volume 51

Caging the Beast: A theory of sensory consciousness
by Paula Droege

Caging the Beast

A theory of sensory consciousness

Paula Droege

Pennsylvania State University, Pennsylvania

LIBRARY
TRENT
UNIVERSITY

John Benjamins Publishing Company

Amsterdam/Philadelphia

B808.9 .D76 2003

™ 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
TRENT
UNIVERSITY

Library of Congress Cataloging-in-Publication Data

Droege, Paula
Caging the beast : a theory of sensory consciousness / Paula Droege.
p. cm. (Advances in Consciousness Research, ISSN 1381–589X ; v. 51)
Includes bibliographical references and index.
1. Consciousness. 2. Sense (Philosophy) 3. Philosophy of mind. I. Title. II. Series.

B808.9.D76 2003
126-dc21 2003045379
ISBN 90 272 5181 9 (Eur.) / 1 58811 390 6 (US) (Hb; alk. paper)
ISBN 90 272 5182 7 (Eur.) / 1 58811 391 4 (US) (Pb; alk. paper)

© 2003 – 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

To Jon

Table of contents

Preface ix

Chapter 1
On sensory consciousness 1
1.1 Caging the beast 7
1.2 Internal sense: A good trap 17
1.3 Second sense: A better trap 21
1.4 Purely verbal? 25

Chapter 2
On higher-order theories of consciousness 31
2.1 The higher-order explanation of state consciousness 32
2.1.A Higher-order thought theory 32
2.1.B Objections to the higher-order thought theory or playing the shell game 38
2.1.C Higher-order perception theory 44
2.1.D Objections to the higher-order perception theory 46
2.2 Formulating an alternative: A flat theory of sensory consciousness 52
2.2.A Dretske's trouble spot 54
2.2.B Spot-sight and thimble-seeking 59
2.3 Conclusion 63

Chapter 3
Solving the problem of Spot-sight 65
3.1 Coordinating sensory consciousness 66
3.1.A The best approximation of the world 73
3.1.B Representing 'now' 75
3.1.C Decision and action 81
3.2 What good is a second sense? 86
3.3 Spot-sight again 95

Chapter 4
Subjectivity 99
4.1 Subjective authority **101**
4.2 Special facts or special access? **105**
4.2.A On what it's like **105**
4.2.B Nagel's funny facts **106**
4.2.C Point of view provides a special route, not special facts **109**
4.3 Subjectivity as the view from here **110**
4.3.A Tokens in a language of thought **110**
4.3.B The reciprocality of subject and object **114**
4.3.C Egocentric maps **116**
4.4 Deflating (and re-inflating) subjectivity **120**

Chapter 5
Testing the theory 123
5.1 Troubles with functionalism **123**
5.1.A Chauvinism **124**
5.1.B Liberalism **125**
5.2 The hard problem **128**
5.3 On Rosenthal **132**
5.4 Dealing with Dennett **135**
5.4.A The Camera Obscura argument **135**
5.4.B The bizarre category of the objectively subjective **138**

Appendix
A Speculative Hypothesis 141
A.1 Attention as the coordination of sensory representations **142**
A.2 Locating sensory consciousness **145**
A.2.A Many senses or one? **146**
A.2.B Looking for the second sense **148**
A.2.C A home for conscious sensory states **152**

Notes 157

References 169

Index 179

Preface

> *Falling asleep has never struck me as a very natural thing to do. There is a surreal trickiness to traversing that in-between area, when the grip of consciousness is slipping but has not quite let go and curious mutated thoughts pass as normal cogitation unless snapped into clear light by a creaking door, one's bed partner twitching, or the prematurely jubilant realization* I'm falling asleep.
>
> John Updike

The converse transition, from sleep to waking, can be just as unsettling. Rousing from a deep slumber is confusing, rendering strange the otherwise familiar sounds of home and family. The bright light of day dawns slowly at these times, beginning with a gray grogginess followed by half-waking fumbling with alarm clock and bathrobe. This transition between waking and sleeping is an excellent *intuitive* analogy for the distinction between unconscious and conscious sensory states that is the topic of this book. When awake we have conscious states: we experience the passage of time by the changes in scenery as we move about the world. When in dreamless sleep we have unconscious states: the world disappears and time seems to pass in an instant. The difference is dramatic and calls for explanation. But it is not a good *analytical* analogy. As I sit here at the computer, awake, I am undergoing all kinds of sensory states that remain unconscious, such as hearing the sound of the computer humming or feeling the cramp in my shoulder. Likewise, some sleep states, such as when dreaming, may count as conscious states.

In keeping with the analytical tradition, this book proposes various refinements of the intuitive difference between waking and sleeping toward an explanation of the phenomenon I call *sensory consciousness*. We exhibit sensory consciousness, I will argue, when we have sensory states that are coordinated into representations of the world at the present moment. Unconscious sensory states, such as those we undergo in dreamless sleep, are not coordinated into a representation of the present moment. As is immediately clear, the theory counts itself among representational theories of consciousness, and I offer it primarily as an addendum to the productive work ongoing in this field. In reading representationalist theories of consciousness, it always seemed that the

phenomenon of sensory consciousness was unexplained by the theory. In some cases (Dreske 1995; Tye 1995, 2000) the theory seemed to offer an explanation of the differences among conscious sensory states – why the representation of a red ball has a different qualitative character from the representation of a green square – yet failed to explain why these representations are conscious at all. In other cases (Rosenthal 1997; Lycan 1995) the theory seemed to offer an explanation of introspection rather than of sensory consciousness. In Chapters 1 and 2 I will develop these worries, and articulate the phenomenon of sensory consciousness to be explained. In the remaining chapters I will offer my own representationalist theory, called the *second sense theory*. Regardless of the fate of the particular theory I propose, I will count the book a success if representationalists recognize the phenomenon of sensory consciousness and see the ways in which it differs from sensory representation and introspective representation.

The primary thanks I owe in the production of this book go to my advisor at the University of Connecticut, Austen Clark, for thoughtful comments on previous drafts and useful advice for negotiating the murky and treacherous territory of consciousness theory. Thanks as well to Ruth Millikan and Crawford Elder for suggesting various clarifications and elaborations. A presentation to the CUNY Cognitive Science Group was enormously helpful in articulating the claims about higher-order thought theory, and more generally I am grateful to the group for providing a stimulating environment for working through questions in philosophy of mind. Similarly, the NEH seminar on Consciousness and Intentionality, led by David Chalmers and David Hoy, presented an unrivaled opportunity for discussion and thought with a collection of engaging, acute, and subtle colleagues. My conversations with two cognitive neuropsychologists, David LaBerge and Sharon Hannigan, have been invaluable in my search for connections between attention theory and consciousness. As much as I have learned about the workings of the brain, I realize there is vastly more to learn, and I look forward to exploring the ways philosophers and psychologists can continue to learn from one another. Finally, I appreciate the advice and comments from the editor of this series, Dr. Maxim Stamenov, and anonymous reviewers. More than anyone, though, my husband Jon and son Noah deserve the credit for bringing this book to fruition. Without their love and support the world would be a far drearier place for me, demonstrating that sensory consciousness is the mere minimum for a good life.

Paula Droege
Rhinebeck, NY

Chapter 1

On sensory consciousness

Consciousness, like love, is something so intimate and vital to our sense of ourselves as human, that explanation, even definition, risks robbing humanity of its soul. The proposal that consciousness could be fully explained in terms of physical stuff like neurons challenges the idea that we are special, mysterious beings, deserving of considerations such as inalienable human rights. Yet we need not think of mystery and explanation as incompatible. Consider life. Despite knowing the basic facts about human reproduction – DNA, cell division and the like – giving birth was the most mysterious process I have ever experienced. The sense of mystery here seems to come from sources other than inexplicability: complexity, cultural significance and symbolic value. Similarly with consciousness, I do not expect the mystery of consciousness to be removed when it has been successfully explained. At minimum a successful explanation must provide criteria for determining when consciousness is present and when it is not. This minimum standard will be my goal in the following. A fuller account would include a description of the mechanics of consciousness and an explanation of why consciousness exists in the first place. Though I will hint at possible hows and whys of consciousness, we will see that my more modest goal presents enough problems on its own.

The first problem is to establish exactly what phenomenon of consciousness is at issue. As a moment's reflection reveals, 'consciousness' refers to a diverse range of phenomena. While not exhaustive or precise, the following list gives a sense of the breadth and variety of uses for the term 'consciousness'. Consciousness can simply mean being awake rather than asleep or in a coma. The phrase 'regained consciousness' is usually meant in this way. Consciousness can mean having sensations and thoughts. When I sense or think about something, I am conscious of it in some sense. But here already we face an additional ambiguity in the words 'sensation' and 'thought'. Take the sensation of pain, for example. When we complain about our pains, we normally are referring to the characteristic hurtfulness of pain. But we also may speak of a pain that continues, like an all-day headache, even though the hurtfulness comes and goes. I will come back to this important ambiguity in later

discussion. For now we can simply point to two senses of 'conscious sensations and thoughts': having sensations and thoughts, or *feeling* one's sensations or thoughts in characteristic ways.[1]

The last two ways of using 'consciousness' that I will mention are, to my mind, significantly different than the previous forms. Consciousness can mean being aware of one's sensations and thoughts, as when one is conscious of being afraid or of thinking inappropriate thoughts. Additionally, consciousness can mean being socially aware, as when one is influenced in some way by the culture, language and social systems in which one lives. These forms of consciousness seem to require a higher level of intellectual and social development. Very lowly beasts undergo some form of sleep/wake cycle, and even plants have sensations in the sense of responding to physical stimuli. We can thus say that beasts and plants are 'conscious' in some sense but could not conclude that they are aware of their sensations or participated in the current *Zeitgeist*. Clearly, we will need to be more specific about what form of consciousness is to be explained in order to determine when 'consciousness' is present and when it is not.

My target will be the sort of consciousness one has when one has sensations and thoughts. Keeping in mind the ambiguity noted above, we can be a bit more specific and say it is the sort of consciousness one has when one *feels* one's sensations and thoughts in characteristic ways. To narrow the field further still, I will deal exclusively with conscious sensations. So let us call the form of consciousness at issue *sensory consciousness* and call the individual sensations exhibiting this form of consciousness *conscious sensory states*. An explanation of sensory consciousness, then, will allow us to distinguish conscious sensory states from unconscious ones. In the example of the all-day headache, an adequate explanation should tell us what factors determine when we feel our pain and when we do not. What is different when the pain exhibits its characteristic hurtfulness, as opposed to when it does not?

To be clear, I do not claim to analyze the meaning of the concept of 'consciousness.' If the philosophical debate about consciousness shows anything, it shows that 'consciousness' does not refer to a unitary phenomenon and so is a poor subject for conceptual analysis. Rather, I aim to identify an empirical phenomenon and to provide an operational definition of its essential features. Once we are reasonably sure we know what phenomenon is at issue, and regarding consciousness this is no mean task, we should throw away the ladder of ordinary language and focus on the explanation of the phenomenon identified. Thus, we should be prepared for the possibility that the resulting theory will postulate a definition of sensory consciousness that does not fit neatly with

common usage. And so we may need to reconsider some of the ways we commonly talk about consciousness. While some have suggested that the ambiguities of 'consciousness' in common usage requires we drop the term entirely, (Wilkes 1988; Churchland 1997) I doubt things will come to such an extreme. More likely we will simply need to revise some of our assumptions about what sensory consciousness is and how we come to possess it.

I should also note how I am using the term 'mental' in identifying a 'mental' phenomenon. I hold a version of Brentano's thesis that the mental just is the intentional[2] including in some sense its functional role. The present proposal follows others who hold this representationalist view in attempting to account for mental phenomena such as consciousness in representational terms. So I will argue that sensory consciousness can be explained entirely in terms of mental representations and their functional organization. In keeping with this view, I take sensory states to have the function of representing features of the world, such as colors, shapes, sounds and textures.[3] The sensations I am now having represent such things as a grey laptop emitting a low hum, a bird chirping as it flies by the window, and the green-gold leaves of a stately tree swaying in the breeze. Though I will not argue for this view here, I refer the skeptic to the sensory representation arguments of Michael Tye (1992, 1995, 1998, 2000), Austen Clark (1993, 2000), William Lycan (1987, 1996) and others.

Similarly, I will not argue for a particular theory of mental representation. I favor a teleo-functional account of the form offered by Ruth Millikan (1984, 1993). On such an account, R is a representation of X if it is produced by a mechanism whose proper function is to make items whose forms vary isomorphically with conditions in the world, according to a rule, and this isomorphic relation has proven useful to the creature in adapting its behaviors to these conditions. As anyone who has read Millikan's work knows, this brief account is a gross oversimplification of a very complex view. Yet it captures the central features of the theory that are particularly salient to my theory of sensory consciousness: the necessity of function in individuating intentional content and the possibility that a representation may fail to fulfill its function while maintaining its intentional content. These elements will be important in distinguishing conscious sensory states from unconscious sensory states and in distinguishing conscious sensory states from hallucinatory or illusory conscious states.

Before we can tackle these issues we need to clarify the notion of 'consciousness' at issue. To begin, we can usefully distinguish between questions about the contents of conscious sensory states and the more general question of what makes those contents *conscious*. Conscious sensory states have all sorts

of contents – the smell of morning coffee, the sound of a Mozart concerto from the stereo, the feeling of exhaustion, you name it. The phenomenon of sensory consciousness I wish to consider is that which is common to all conscious sensory states, regardless of their content. Much of the argument about consciousness since the revival of the problem two decades ago has revolved around the contents of conscious sensory states, in particular the phenomenal content known as 'qualia'. Unfortunately, 'qualia' is used as ambiguously as 'consciousness' and 'sensation', so it is now nearly useless as a technical term. For present purposes, I will use the term *qualitative character* to mean that which determines the characteristic way something feels. When I look at grass, why does it look green to me rather than red? Might qualitative character be inverted, so things that look green to me look red to you? How does the way round things feel differ from the way round things look? These questions about qualitative character are some of the most puzzling problems in philosophy of mind. But these questions, as with any questions about a particular sort of content, can be separated from the question of sensory consciousness per se. Even presuming we could solve problems of qualitative character and other sorts of mental state content, there is still the question of what makes those mental states *conscious* as opposed to unconscious.[4]

In his review of consciousness theory, Joseph Levine puts the point this way:

> There are two questions a Materialist theory has to answer: (1) what distinguishes conscious experiences from mental states (or any states, for that matter) that aren't experiences? and (2) what distinguishes conscious experiences from each other; or, what makes reddish different from greenish, what determines qualitative content? (Levine 1997: 388)

The first of these questions, call it the *question of state consciousness*, is the question I propose to answer in part. In other words, conscious sensory states form a subset of conscious states, so the an answer to the question of sensory consciousness constitutes a partial answer to the question of state consciousness. The second question, call it the *question of qualitative character*, can be answered separately from the first because sensory states can be distinguished by what features of the world they represent.[5] In other words, sensory states can be individuated in terms of intentional content. We can call a sensory state 'red' (or 'reddish') as opposed to 'green' (or 'greenish') just in case it represents red features as opposed to green features. Now, how we go about determining what counts as a 'red' feature and whether there is more to 'red' than its representational content are fraught questions, as witnessed by the expanding literature

on color and color perception. My point here is that the question of qualitative character is separate from the question of state consciousness because both unconscious and conscious sensory states can have qualitative character. In the case of the all-day headache, whatever features of the sensory state distinguish a throbbing headache from a stabbing headache arguably applies when that state is not conscious. Discrimination tests as in blindsight cases and priming studies suggest that both unconscious and conscious sensory states can represent sensory features like color. Blindsight patients, for example, deny they are having any conscious sensory experience yet are able to guess with better than chance accuracy as to whether the stimulus presented in the blind area of their visual field is red or green. (Farah 1997; Stoerig 1998; Stoerig & Cowey 1996) These results suggest that whatever distinguishes red sensory states from green sensory states is present even when the state is not conscious.[6]

The distinction between the question of qualitative character and the question of state consciousness is an important one in the sort of 'divide and conquer' strategy for explaining consciousness recommended by William Lycan (1996:2). When this distinction is blurred, theories about qualitative character are easily misread to be theories about state consciousness, especially since theories of qualitative character often claim to be theories of 'consciousness.' Take, for example the notion of 'phenomenal consciousness' described by Michael Tye and Ned Block. On first look, phenomenal consciousness seems identical to my target explanandum of sensory consciousness. Tye describes phenomenal consciousness as associated with "talk of 'raw feels' and 'what it is like.' For a person who feels pain, there is something it is *like* for him to be in pain. *Phenomenal consciousness* is present." (Tye 2000:ix) Similarly, Block identifies the phenomenal properties of a state as those that make up "what it is like" to have it. They are the "experiential properties of sensations, feelings, and perceptions. . .thoughts, desires, and emotions." (Block 1995:230)

Yet both hold that phenomenal consciousness exists when there is no state consciousness. Block describes a case similar to the all-day headache:

> suppose you are engaged in intense conversation when suddenly at noon you realize that right outside your window there is – and has been for some time – a deafening pneumatic drill digging up the street. You were aware of the noise all along, but only at noon are you *consciously* aware of it. (Block 1995:234)

Block presents this as a case where you were phenomenally conscious all along. The difference is that at noon you also acquired what he calls 'access consciousness.' Access consciousness may be Block's answer to the question of state consciousness, but I will argue in the next section that it is insufficient. The point

here is that phenomenal consciousness clearly is an answer to the question of qualitative character. And once we have pulled apart qualitative character from sensory consciousness, it is no longer obvious that there is 'something it is like' to be in a state of phenomenal consciousness, or that phenomenal consciousness is the sort of consciousness that seems a scientific mystery. (Block 1995:230) To push the point even further, if a person has a sensation of pain while in dreamless sleep, then she is in a state of phenomenal consciousness. (Block 1995:235) Here it seems reasonable to say that this person is having a sensation with the character of pain, but she does not feel the characteristic hurtfulness associated with the pain. The pain is an unconscious state.

Tye treats the same cases in a similar way. Of the ongoing jackhammer experience he says: "In one sense, my experience was unconscious: I was not aware that I was having that experience. This is a failure in higher-order or introspective consciousness. But it does not follow from this that there was no phenomenal consciousness." (Tye 1997:294) Of pains when in a state of dreamless sleep, Tye would also say they are cases of phenomenal consciousness.[7] Again, the point is that Tye's theory of phenomenal consciousness is an answer to the question of qualitative character, not to the question of state consciousness *per se*.

My emphasis on this distinction should not be read as an argument for some theoretical hierarchy where the question of state consciousness is the *real* question of consciousness or a more important or difficult question to answer. On the contrary, I believe the progress in answering these two questions is mutually reinforcing. As we learn more about what accounts for or might account for differences in qualitative character, the way those differences are exhibited in conscious states becomes less mysterious. Likewise, gaining a better sense of the distinction between conscious and unconscious sensation may help alleviate the impression that qualities such as 'redness' have an eerie metaphysical status incompatible with physical systems. At least, one can hope. In any case, I make no claim to be solving *the* problem or *the* mystery of consciousness. There are plenty of mysteries to go around.

My question, to reiterate, is what distinguishes unconscious sensory states from conscious ones? What distinguishes the sensory state of a blindsight patient representing a red stimulus from the sensory state of a normally sighted person representing a red stimulus? As I have said, my discussion is limited to sensory states, such as when one visually represents a red or green stimulus.[8] So the state you are in now of seeing the words on the page or feeling the paper turn in your hand are the sort of mental states I will be considering. I will not argue that the theory applies to conscious thoughts, although I think

it could be extended to cover them. The two main reasons for this choice are: (1) Conscious sensory states are arguably more basic than conscious thoughts. Indeed, I will argue that there can be, logically if not practically, conscious sensory states without any concepts whatsoever. (2) As the discussion of qualitative character suggests, sensory states have been the primary focus of discussion about consciousness generally, so a theory of consciousness should account for at least this sort of mental state. I firmly believe that progress in consciousness theory can only be made by carefully subdividing the multifarious forms of phenomena that the word picks out and considering each in turn. A later synthesis of the theories that result from this process is bound to be more effective than the current frustrating search for a single account of all the phenomena of consciousness so-called.

1. Caging the beast

In order to explain sensory consciousness, we will need to get clearer about what conscious sensory states are. This will not be an easy task. There is so little agreement about anything regarding consciousness that one philosopher jokingly claimed the persistent disagreement about definitions to be an existence proof for the ineffability of consciousness.[9] Because so many different terms are used to describe the various phenomena involved, it can be difficult to determine when the dispute is purely verbal and when substantive disagreement lies behind terminological differences. To offset these terminological problems somewhat, I begin with very simple taxonomy of three kinds of sensory state: unconscious, conscious and self-conscious. As I describe these kinds of sensory state I will consider how they interrelate and how other descriptions of consciousness fit into this schema. Though I restrict my explanandum to conscious sensory states, not everyone does the same. So I will use the terminology and descriptions of the authors where appropriate, drawing connections to my target, conscious sensory states, where applicable.

> *Unconscious sensory state.* A creature has an unconscious sensory state when the state is a sensory representation but there is 'nothing it is like' (Nagel 1974/1991) to have that sensory state.[10]

The idea of unconscious mental states is relatively new. For Descartes and Locke, who inaugurated the discussion of consciousness, all mental states were conscious states. Consciousness characterized mentality. Descartes maintained that "there can be nothing in the mind, in so far as it is a thinking thing, of

which it is not aware." (Descartes 1641/1984:171)[11] The same idea appears in Locke where he says, "thinking consists in being conscious that one thinks." (Locke 1689/1975:II, I, §19, 115) But already in these early days of consciousness theory, it is unclear what sense of 'conscious' is being used. Descartes supposed that a person is always thinking and therefore always conscious. (Descartes 1641/1984:18) To Locke, the notion that a person is conscious while sleeping was ridiculous. (Locke 1689/1975:II, 1, §11, 110) This disagreement about the relation between sleeping and consciousness suggests that Descartes and Locke were using different senses of either 'thought' or 'conscious' or both. While both claimed that all mental states were conscious states, they seem to have very different notions of what this meant.

As philosophers have developed theories of consciousness, the idea that all mental states are conscious has lost favor. David Rosenthal in particular has argued persuasively for the possibility of unconscious mental states, noting that there are many cases where a person's behavior or mental activity is best explained by the existence of unconscious mental states. The all-day headache, as I have noted, is reasonably considered to be a single sensation that is intermittently conscious. One does not feel the pain when distracted – there is nothing it is like to feel the pain – so during that time of distraction the sensory state is unconscious. Yet it is reasonable to say the same headache endured throughout. As Rosenthal observes, "it would be odd to say that one had had a sequence of brief, distinct, but qualitatively identical pains or aches." (Rosenthal 1997:731) Rosenthal also points to the way we attribute thoughts, desires or emotions to a person even when those mental states are not at all conscious. (Rosenthal 1997:731) I may be angry about something but not be aware of my anger until someone else remarks on my scowls and quick temper. Freudian repressed states explain otherwise random neurotic behaviors, and unconscious reasoning processes explain the 'light bulb phenomenon', as when the answer to a problem suddenly occurs to you even though you had been consciously thinking about something else.

Current psychological literature provides additional examples of unconscious sensory states. As noted earlier, in blindsight studies patients are asked to make guesses about objects presented in the blind area of their visual field. Because the guesses of blindsight patients are accurate at a rate better than chance, psychologists conclude that visual stimuli from the blind area are processed unconsciously. (Stoerig 1998) There is 'nothing it is like' for the blindsight patient to have these mental states. More familiarly, conscious processes such as speech are usually subserved by unconscious sub-processes (parsing, word choice, etc.) that are routine and operate in parallel. (Dennett 1991)

Much of the psychological research on implicit or covert processing may be taken as an investigation of unconscious sensory states. In this sort of research, psychologists test the influence of stimuli presented 'sub-threshold', where a 'threshold' for sensory consciousness is determined by the conditions under which a subject is able to report a stimulus. When a stimulus is presented under sub-threshold conditions, subjects report no stimulus.[12] Semantic priming, for example, involves briefly presenting a word followed by a masking word. Though subjects do not report seeing the first word, studies demonstrate its effects on responses to other words. The brief, sub-threshold presentations can be interpreted as 'unconscious' (Marcel 1983). Since the threshold of sensory consciousness is determined by reportability of the stimulus, however, there is some question whether 'sub-threshold' truly indicates an unconscious sensory state.[13] Nonetheless, these cases show the need for a distinction between unconscious and conscious sensory states in psychological research and suggest possible examples of just this distinction.

Second, Rosenthal argues that if all mental states are conscious, no explanation of consciousness is possible. We cannot explain consciousness in terms of the mental if consciousness is built into the nature of mentality in the way Descartes and Locke suggest. The only alternative form of explanation, one in purely physical terms, is a daunting project, if not an impossible one. As Rosenthal states (perhaps a bit too strongly), "nothing that is not mental can help to explain consciousness. So, if consciousness is essential to mentality, no informative, nontrivial explanation of consciousness is possible at all." (Rosenthal 1991c: 463) We would have no other choice but to accept consciousness as primitive.[14]

Despite such arguments, remnants of the identification of mentality and consciousness survive, even in Rosenthal's own theory. In "A Theory of Consciousness" (1997), Rosenthal distinguishes between two uses of the word 'conscious', transitive and intransitive consciousness. "One is when we speak of our being conscious of something. Because of the direct object, I shall call this the transitive use." (Rosenthal 1997: 737) Taken as simply a grammatical locution, 'conscious of' is a useful and unobjectionable phrase, reasonably distinguished from the use of 'conscious' as applied to mental states. When we say a mental state is 'conscious' there need be no direct object, so Rosenthal suggests we call this second usage the 'intransitive' use. (Rosenthal 1997: 737) But there is more than a grammatical distinction involved here. Rosenthal needs there to be an ontological distinction as well because he means to "explain intransitive consciousness in terms of transitive consciousness." (Rosenthal 1997: 737) If tran-

sitive and intransitive consciousness are not distinct, the explanation would be circular. Thus, Rosenthal notes that

> transitive consciousness can occur without intransitive states consciousness. 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. And if, as is likely, mental states are possible during sleep, transitive consciousness will not even presuppose creature consciousness.[15] (Rosenthal 1997:737)

Elsewhere Rosenthal writes: "It is obvious that when we have thoughts about things, we are transitively conscious of those things. ... So being in a mental state is very often sufficient for one to be transitively conscious of something." (Rosenthal 1993b:356)[16] These descriptions suggest that a mental state with intentional content, a state that is of or about something, can suffice for one to be transitively conscious of something. In other words, having a mental state that is 'of or about' something can be sufficient for a person to be 'conscious of' that thing, even if the mental state is repressed or otherwise an unconscious sensory state of the sort described in the previous paragraphs. Moreover, an unconscious person (a person lacking 'creature consciousness') can be 'conscious of' something. I find it confusing to talk of someone being 'conscious of' something when that person has no conscious states and may herself be unconscious as well.[17] Rosenthal is careful to note that transitive consciousness is not the sort of consciousness requiring explanation, so there need be no particular mystery about how an unconscious person can be conscious of something.[18] Nonetheless, I will argue in Chapter 2 that using the word 'conscious' to refer to persons without conscious states lends unwarranted plausibility to Rosenthal's explanation. As we shall see, substantive consequences can follow from terminological choices.

Certainly people can use terms any way they like, so long as the definitions are clear. In the case of consciousness, however, there is an argument to be made for restricting one's use of the term whenever possible. When we multiply kinds of consciousness we multiply the risk of confusing the phenomena identified by one with the phenomena identified by the other. We must accept this risk in order to distinguish different aspects of the explanandum; such is the 'divide and conquer' strategy mentioned earlier. On the other hand, if there is a term that can adequately identify a referent without using the word 'consciousness', this alternate term would be preferable in order to minimize the already inevitable confusion engendered by the nature of consciousness. So, instead of 'transitive consciousness' I will simply talk of persons with 'intentional'

or 'representational' states.[19] Some other kinds of 'consciousness' that arguably admit of a similar substitution are David Armstrong's 'minimal consciousness' and 'perceptual consciousness' as well as Ned Block's 'access consciousness.'

The first case, Armstrong's 'minimal consciousness', involves only rudimentary mental processing. In contrast to a state of total unconsciousness, such as deep, dreamless sleep, minimal consciousness involves some mental activity, a mental state producing mental effects. Something mental is occurring. Armstrong gives the example of waking up and knowing the answer to some puzzle, what I referred to earlier as the 'light bulb phenomenon' (Armstrong 1981: 56–58; Armstrong & Malcolm 1984: 119). Though a person is asleep, her ongoing mental processing produces the needed answer. Rather than call this minimal consciousness, however, I will simply call this a case of unconscious mental activity. The distinction Armstrong is drawing between minimal consciousness and total unconsciousness is reasonably understood as the distinction between active and inactive mental states. Deep, dreamless sleep is an example of total unconsciousness because the person's mental states, her beliefs, memories and desires, are inactive. She is like a computer that has programs and data storage but is turned off and so is not undergoing any processing activity. The minimally conscious person is doing some internal processing – she is 'on' in some sense – but she has no conscious sensory states. Rather than use a separate term, 'minimal consciousness', I will mark Armstrong's difference between total unconsciousness and minimal consciousness, where relevant, by pointing out the element of mental activity involved.

The second sort of state Armstrong describes is 'perceptual consciousness.' In this sort of state a person is perceiving the environment, but may not be aware of what she is perceiving. Armstrong gives the example of a sleepwalker who manages to navigate through doors, down stairs and around obstructions without being aware of her perceptions (Armstrong & Malcolm 1984: 119). Here again I will refrain from using the term 'consciousness' and will refer to the sleepwalker's ability as 'perception' or more generally as 'sensory representation.' Sometimes sensory representations are conscious sensory states and sometimes they are not. What makes a sensory representation conscious is the principal question I am investigating.

Finally, and most contentiously, Ned Block's 'access consciousness' seems to identify unconscious states. At least, I would argue that the conditions for access consciousness are insufficient to constitute state consciousness. Here are the three conditions Block gives as jointly sufficient but not all necessary for access-consciousness:

> A state is access-conscious (A-conscious) if, in virtue of one's having the state, a representation of its content is (1) inferentially promiscuous ..., that is, poised for use as a premise in reasoning, (2) poised for rational[20] control of action, and (3) poised for rational control of speech. (Block 1995:231)

The key term in Block's definition is the word 'poised'. If 'poised' simply means something like 'available', then access-consciousness is very much like Armstrong's description of total unconsciousness. Access-conscious states would be like the data in a computer that are at present inactive but could at any moment be activated. Block admits that A-consciousness is dispositional, but claims that it is not "totally dispositional" because "then quiescent or inactive beliefs will count as A[-consciousness]." (Block 1997:160) So, 'poised' must have more of an active sense than simply 'available.' Block notes the condition of 'inferential promiscuity' which he clarifies as "free use of a representation as a premise in reasoning." (Block 1997:160). To be poised, the content of a state must be primed and ready for use. Perhaps we could say that access-conscious states are poised for reasoning in the way a dancer is poised to leap. While not yet actually leaping, perhaps not moving at all, the dancer is on stage and ready, anticipating the next move, positioned to act. But here again we would have to ask what it means to be 'on stage' and 'anticipating the next move' in order to flesh out the more active sense of 'poised' Block has in mind.

Block's example of pure access-consciousness is the hypothetical superblindsighter. The superblindsighter is a person who, like blindsight patients, is blind to visual stimuli yet, unlike blindsight patients, has learned how to prompt himself to guess what stimuli are in the blind field.

> The superblindsighter spontaneously says, "Now I know there is a horizontal line in my blind field even though I don't actually see it." Visual information from his blind field simply pops into his thoughts in the way that solutions to problems we've been worrying about pop into our thoughts ...
>
> (Block 1995:233)

Unlike the blindsighter who cannot prompt himself to guess about his own states (a critical functional difference between blindsight patients and normally sighted persons), the superblindsighter has access to information from his visual system. Before it 'pops' into his thoughts, the visual information is available to him as one among many bits of information he can choose to act upon. This way of talking about accessing information is reminiscent of the psychologist's concept of working memory. In working on a problem we may keep several ideas in play by thinking first about one and then another, juxtaposing them in different ways in hope of hitting on a solution. While this sense of

a 'poised' state seems more suited to be called 'conscious' than the first sense of stored data, there is negligible difference in the two kinds of states. In both cases the poised state is not yet active. If the state is actually used in reasoning, action or speech, there may be reason to consider it a form of conscious state. In this case an access-conscious state might also be a conscious state, but its being access-conscious still would not constitute its being a conscious state.[21]

> *Conscious sensory state.* A creature has a conscious sensory state when the state is a sensory representation of something in the world, and there is something it is like to have that state.

As a first stab, this description of sensory consciousness tells us two things: a conscious sensory state is a representational state whose object is external,[22] and there is 'something it is like' to have a conscious sensory state. Neither feature is particularly helpful, however; the first is too broad and the second is too vague. The next two chapters will slowly work toward an operational definition of sensory consciousness through examples and refinements on these conditions. Finally in Chapter 3 I will propose a full definition of conscious sensory states as it fits within a theoretical explanation of sensory consciousness.

To start the definition process, here are some examples:

> Riding my bicycle down a steep hill, I unexpectedly hit a deep groove in the road and am thrown from the bike. I get up, put on my glasses and walk over to pick up my bike. At least, I am later told by witnesses that I have done this. I remember nothing from the time I was riding down the hill until much later being questioned by medics.[23]
>
> For the last 20 minutes I have been shifting around on my chair, crossing one leg and then the other, sitting forward, then back. Only when I turn my attention away from my task do I become aware of the bodily condition that has been causing my movement. Now I notice the ache in my shoulders and the cramped feeling in my legs.

The first case highlights some of the typical differences in mental states that have been taken as markers of consciousness. As noted above, sensory responsiveness is sometimes considered a form of consciousness. My movements to get up, pick up my glasses and bicycle show that I have seen the glasses, bicycle, etc. in some sense, because I have reacted appropriately. But the sense of 'seen' involved in sensory responsiveness is neither a necessary nor a sufficient condition for sensory consciousness. It is unnecessary because motor functions may be disconnected from other mental processes. A person may have fully conscious auditory sensations but be unable to respond to them due to paralysis

or other motor disorder. Therefore, lack of responsiveness does not imply lack of sensory consciousness. Nor is sensory responsiveness sufficient for sensory consciousness. The sort of chair shifting behavior noted in the second example is a case of sensory responsiveness without sensory consciousness. Though the pain from my shoulders and legs was not conscious, I responded to it by adjusting my position several times. Similarly, a person may be in a state of dreamless sleep and move in response to a muscle cramp or a poke from a bed partner. So the sensory representations that facilitate motor responses may not be conscious sensory states.

While there may be some question about whether or not I had conscious sensory states when picking up my glasses and bicycle, it seems more likely that I had conscious sensory states later while I was talking to the medic. Because we tend to remember what we experience consciously, memory has also been taken as an indicator of sensory consciousness. If I remember an event, such as talking to the medics, this memory indicates that I was having conscious sensory states at the time of the conversation. Memories are notoriously fallible, however, and so can only serve as an indicator that the remembered sensory state was conscious. One may be tempted to make a strong claim and say that in the case of *episodic memory* the remembered sensory state *must* have been conscious. In episodic memory a person remembers an event as if she were there at the time as a participant or observer rather than having learned of the event is some other way. I remember looking into the face of the medic, and can recall sensory details, such as the mole on his left cheek, by producing a mental image of the event. But even episodic memories can be distorted if the memory is inconsistent with other memories or beliefs. (Dennett 1991) Perhaps the mole was on the left cheek of the store clerk I had seen moments before the accident and I somehow confabulated the image of the medic. So, even though episodic memory often indicates conscious sensory states at the time of the remembered event, it does not do so reliably.[24]

Episodic memories are not a sufficient criterion for identifying past conscious sensory states, nor are memories in general a sufficient criterion for identifying present conscious sensory states. In other words, memories of all kinds can be unconscious. Right now, presumably, we all possess multitudes of memories that are not conscious. If prompted I can remember what I had for breakfast or who phoned me yesterday. If the event made a sufficient impression, I may even recall some episodic memories, full of sensory detail. But without such a prompt, these memories would likely remain unconscious sensory states. The existence of unconscious memories is evident by their quick and (relatively) consistent recall as well as the way they can influence behavior.

After a bad meal in a restaurant I may avoid returning there, even if I don't consciously remember the event as my reason. Freudian repressed memories have been credited with all sorts of influence, from psychosomatic pains to the formation of multiple personalities. So it is possible to have memories that are not themselves conscious and do not imply the existence of past conscious sensory states.

It is also possible to have conscious sensory states without any episodic memories whatsoever. In cases of anterograde amnesia, patients are unable to form new memories. Though alert and attentive, patients with brain disorders such as Korsakoff's Syndrome may lack normal memory function. Like the case of motor dysfunction described above, these amnesia patients likely have fully conscious sensory states but lack the ability to remember them (Kolb & Whishaw 1990: 552). Though sensory consciousness without episodic memory is possible, some form of memory may be necessary for sensory consciousness. Skill-based memories such as recognition and conditioned response are retained in even the most severe amnesia cases. However, because skill-based memory is so basic to mental function – arguably necessary for perception and the acquisition of simple concepts – its relation to sensory consciousness is not likely to be very informative. The role such memories play in sensory consciousness is probably duplicated for many other mental functions. Only episodic memory seems to be bound in an interesting way to sensory consciousness, but as I have argued, even this form of memory fails to be necessary or sufficient for having conscious sensory states.

Putting aside memory, for now at least,[25] we are left with the vague idea that there is 'something it is like' to have a conscious sensory state. This condition does not get us terribly far in itself, because there is no way to further analyze 'what it's like'. The phrase is used more like a demonstrative than a description, to point to something that cannot be otherwise identified. While it is not always clear exactly what is being pointed out, looking at how different authors use the phrase may help determine some identifying markers for sensory consciousness.

One use for the 'what it's like' locution is in comparing Rosenthal's 'state consciousness' with my 'sensory consciousness'. Rosenthal describes 'state consciousness' as a property of mental states that it is like something to be in, while there is nothing it is like to be in an unconscious state (Rosenthal 1993b: 357).[26] Furthermore, I fully agree with Rosenthal when he says that the question at hand is "what it is for a mental state to be conscious. Assuming that not all mental states are conscious, we want to know how the conscious ones differ from those which are not" (Rosenthal 1997: 729). I restrict my claims to con-

scious sensory states rather than attempting to explain all forms of conscious state, but we both agree that there is 'something it is like' to be in a conscious state.[27] Additionally, conscious states represent something in the world on both our accounts. Though we disagree on what constitutes sensory consciousness, we agree that conscious states represent external objects.[28] Given these two features, I take it that 'sensory consciousness' is a subset of Rosenthal's 'state consciousness'.

Self-conscious sensory states, which comprise my third category, differ from conscious sensory states in their representational object. While conscious sensory states are mental states representing external objects, self-conscious sensory states are mental states representing other sensory states. Yet both are 'conscious' in the sense that there is something it is like to have that state. Thus,

> *Self-conscious sensory state.* A creature has a self-conscious sensory state when the state is a mental state about one's own sensory representations, and there is something it is like to have the self-conscious sensory state.

This description of a self-conscious sensory state is similar to Armstrong's definition of 'introspective consciousness' as "a mental event having as its (intentional) object other mental happenings that form part of the same mind" (Armstrong & Malcolm 1984:108). A reconstrual in representational terms might be: one's self-conscious sensory states represent one's own sensory representations. I use the term 'self-conscious' rather than 'introspective' to emphasize the continuity between conscious sensory states and self-conscious sensory states. In my view, self-conscious sensory states are very much like conscious sensory states, except with respect to their representational object. While garden-variety conscious sensory states represent the world, self-conscious sensory states represent one's own sensory states. Since my target is exclusively conscious states, I will not elaborate on or argue for this description of self-conscious sensory states except to mark a distinction between conscious states and self-conscious states. The nature of self-conscious states is the source of considerable disagreement, so I will focus on the distinction in representational object between self-conscious states and conscious states in what follows. To help keep this distinction clear, I will restrict my use of the term 'conscious state' to refer to mental states that represent the world[29] and will use 'self-conscious state' to refer to mental states that represent one's mental states. The question of sensory consciousness, as I construe it, is the question of what determines whether a sensory state is a conscious state or whether it is an unconscious state.

2. Internal sense: A good trap

Now that we have a general view of what sort of beast a conscious sensory state is, what sort of trap might capture it? The one I favor is a theory inspired by inner sense accounts of sensory consciousness. Two key elements figure in inner sense theories: (1) The operation of some sort of sense explains the difference between unconscious sensory states and conscious sensory states. (2) The sense is 'inner' in that it produces representations of internal, sensory states rather than representations of external objects. Of these two elements, my proposal incorporates only the first. Sensory states are conscious when scanned by some kind of sense, but the representations produced are not representations of the sensory states, on my account. So the theory I propose is not an 'inner' sense theory. Nonetheless it is interesting to trace the notion of an 'inner sense' to see the ways it evolved over the course of its history. This project is worthwhile for two reasons: it sheds light on the similarities and differences between an inner sense account and the account I propose, and it attributes many features to the inner sense that many contemporary theories attribute to consciousness, such as enhancing discrimination and coordination capacity as well as making information available for decision-making prior to action. In other words, the idea of some sort of mechanism to fulfill the functions we now attribute to consciousness is an idea that has been around for quite some time.

Aristotle may provide the deepest root of the term. In the section "Common Sense" of *De Anima*, Aristotle notes that each sense discriminates differences among the qualities it senses. Sight discriminates white and black, taste discriminates sweet and bitter. Continuing, Aristotle asks:

> Since we also discriminate white from sweet, and indeed each sensible quality from every other, with what do we perceive that they are different? It must be by sense; for what is before us is sensible objects. (III, 2, 426b10)

Yet Aristotle opposes the idea of a sixth sense in addition to the other five (III, 1, 424b20), and instead argues for a puzzling mereological relation whereby the senses are both many and one. He writes: "In one sense, it is what is divided that perceives two separate objects at once, but in another sense it does so *qua* undivided; for it [the common sense] is divisible in its being, but spatially and numerically undivided." (III, 2, 427a1)

In an effort to clarify the relation between the external senses and the 'common sense', Avicenna (Ibn Sina) proposes a division of mental labor. Additional to the five external senses, Avicenna describes five internal senses: the *sensus*

communis, retentive imagination, memory, estimative faculty and compositive imagination.[30] Avicenna explains:

> It is the function of certain internal faculties to combine certain perceived forms and intentions with others and to separate some of them from others, so that they perceive and also act on what they have perceived. Perception unaccompanied by action takes place when the form or the intention is merely imprinted on the sense organ without the percipient having any power to act upon it at all. (Avicenna, 10th c./1952:30)

For Avicenna, the *sensus communis* has the particular function of coordinating the forms transmitted from the five external senses. The other internal senses retain, combine and divide images.[31] In his commentary on Avicenna, F. Rahman notes that, while the term 'internal sense' is not Aristotelian, all of the functions of the internal senses are found in Aristotle (Rahman 1952:77).

St. Thomas Aquinas reduced Avicenna's internal senses from five to four, but maintained the idea of a *sensus communis* or 'common root of the senses.' "Particular senses discern the particular sense-stimuli proper to them, but to distinguish white from sweet we need some common root sensitivity in which all sense-perceptions meet, and where we can perceive perception itself and become aware that we see." (Aquinas 1273/1989:II, 78, 4) Here Aquinas gives the *sensus communis* the new role of introspection. Not only does the *sensus communis* unite and compare sensory qualities, it is by virtue of its operations that we are *aware* that we see. According to Aquinas, self-knowledge is a matter of knowing the activity of our minds, by virtue of perceiving that activity through the *sensus communis* (II, 87, 1–3).

In John Locke's use of the term 'internal sense,' introspection becomes its exclusive function. Locke presents the internal sense as an epistemological device, a route to knowledge about mental operations in the way the external senses are a route to knowledge about the world.

> The other Fountain, from which Experience furnisheth the Understanding with *Ideas*, is the *Perception of the Operations of our own Minds* within us, as it is employ'd about the *Ideas* it has got. … This source of *Ideas*, every Man has wholly in himself: And though it be not Sense, as having nothing to do with external Objects; yet it is very like it, and might properly enough be call'd internal Sense. (Locke 1689/1975:II, 1, §4, 104)

To avoid confusion with sensation, the first fountain of knowledge, Locke called this second fount 'reflection': "that notice which the Mind takes of its own Operations, and the manner of them, by reason whereof, there come to be *Ideas* of these Operations in the Understanding" (Locke 1689/1975:II, 1,

§4, 104). Examples of the sorts of ideas produced by reflection are perceiving, thinking, doubting, believing, knowing and willing. Sensation, on the other hand, produces ideas about external objects such as yellow, heat, soft and sweet (Locke 1689/1975: II, 1, §3, 104).

Locke, one of the first consciousness theorists, complicated the notion of an internal sense further by identifying it with consciousness. As noted earlier, Locke opposed the Cartesian idea that a person is always thinking. In particular, Locke believed it is obvious that a sleeping person does not think. If thought were possible during sleep, Locke argued, then a person could be happy or miserable without being conscious of it. But this seemed to Locke to be "utterly inconsistent and impossible." (Locke 1689/1975: II, 1, §11, 110) He concluded that thinking requires that we are conscious of thinking, so consciousness must be a matter of reflection. "Whereas hunger consists in that very sensation, as thinking consists in being conscious that one thinks. If they say, That a Man is always conscious to himself of thinking; I ask, How they know it? Consciousness is the perception of what passes in a Man's own mind" (Locke 1689/1975: II, 1, §19, 115).

Always ready to turn the complicated into something truly mind-bending, Immanuel Kant adds the dimension of time to internal sense. Kant begins with an epistemic division between external and internal senses similar to Locke's. Through the external senses we represent objects in the world; through the internal sense we represent states of the mind. According to Kant, each sense structures its objects in a particular form. Whereas the external senses relate objects to one another in space, the internal sense orders internal states according to time. The structure of space and time provides unity and order to objects, making perception possible.

> By means of the external sense (a property of the mind), we represent to ourselves objects as without us, and these all in space. Therein alone are their shape, dimensions, and relations to each other determined or determinable. The internal sense, by means of which the mind contemplates itself or its internal state, gives, indeed, no intuition of the soul as an object; yet there is nevertheless a determinate form, under which alone the contemplation of our internal state is possible, so that all which relates to the inward determinations of the mind is represented in relations of time. (Kant 1781/1990: 23, A23/B37)

As with Locke, the object of the internal sense is one's own mental states. It is worth noting that mental states produced by the external senses are included among the objects of Kant's internal sense. Thus our representations of external things are constrained by the form of time as well as the form of space.

> [B]ecause all representations, whether they have or have not external things for their objects, still in themselves, as determinations of the mind, belong to our internal state; and because this internal state is subject to the formal condition of the internal intuition, that is, to time – time is a condition *a priori* of all phenomena whatsoever – the *immediate* condition of all internal, and thereby the *mediate* condition of all external phenomena.
>
> (Kant 1781/1990:30, A34/B50)

By making sensory representations (those produced by external senses) the objects of an inner sense, Kant establishes the relation between the internal sense and sensory representations that is central to contemporary inner sense explanations of state consciousness. In this explanation, which adds yet another dimension to the notion of an internal sense, mental states become conscious by being the objects of an 'inner sense.'

The inner sense theory made its appearance as an explanation for state consciousness in David Armstrong's *Materialist Theory of Mind*. Citing Kant, Armstrong describes the inner sense as an "awareness of our own mental states" and compares inner sensing to external sensing. "By sense-perception we become aware of current physical happenings in our environment and our body. By inner sense we become aware of current happenings in our own mind" (Armstrong 1968/1993:95). Armstrong calls the awareness that results from inner sensing 'introspective awareness' or in later work 'introspective consciousness.' It is this sense of consciousness, Armstrong argues, that calls for explanation. "Introspective consciousness seems like a light switched on, which illuminates utter darkness. It has seemed to many that with consciousness in this sense, a wholly new thing enters the universe" (Armstrong 1981:63). Armstrong's famous example of a person who is *lacking* introspective consciousness is the long-distance driver:

> This is something that can happen when one is driving very long distances in monotonous conditions. One can 'come to' at some point and realize that one has driven many miles without consciousness of the driving, or, perhaps, anything else. One has kept the car on the road, changed gears, even, or used the brake, but all in a state of 'automatism.' (Armstrong 1968/1993:93)

When he 'comes to' the driver once again becomes aware of his mental states; that is, he regains introspective consciousness.

Far from being a 'wholly new thing', however, Armstrong believes that introspective consciousness is no more mysterious than sense perception. In the same way our external senses are 'directed' at external objects, our inner sense is 'directed' at our mental states (Armstrong 1968/1993:94). One good way to

think of the 'directedness' of perception is in terms of representation. Our external senses produce representations of external objects and our inner sense produces representations of our mental states. As William Lycan puts it, introspective consciousness is "a perceptionlike second-order representing of our own psychological states and events" (Lycan 1996: 13).

What is wholly new about the contemporary inner sense theory is the claim that introspective consciousness can explain state consciousness. As with Rosenthal, Lycan talks generally about conscious states rather than restricting himself, as I have done, to conscious sensory states. So here again I take sensory consciousness to be a subset of state consciousness. According to Lycan, what distinguishes conscious states from unconscious states is just the second-order representation of introspective consciousness (Lycan 1996: 13). When an inner sense scans a mental state, its output is a representation of that mental state. A mental state is conscious when it is represented in this way. "The inner sense theory has it that conscious awareness is the successful operation of an internal scanner or monitor that outputs second-order representations of first-order psychological states" (Lycan 1996: 31). So, for example, the mental states of the long-distance driver again become conscious – he 'comes to' – when his inner sense resumes production of representations about them.

The Armstrong/Lycan version of the inner sense theory is known as a 'higher-order' theory because it explains state consciousness in terms of higher-order representation of mental states. As I said, I reject the higher-order structure of inner sense theories; I will go into more detail about the problems with higher-order theory in the next chapter. In the remainder of this chapter I want to focus on the sensory aspect of an 'internal sense' as a promising feature for an explanation of sensory consciousness.

3. Second sense: A better trap

Attracted by the inner sense theory, I began to study the Armstrong/Lycan explanation of state consciousness, only to be plagued by a certain nagging doubt. The doubt was whether the higher-order inner sense theory does explain state consciousness or whether the theory actually explains self-consciousness instead. Already the terminology of inner sensing as 'introspective consciousness' suggests a mind-mind relation rather than the mind-world relation of state consciousness. Indeed, introspective consciousness is defined as second-order representing, mental states about mental states. So how does the theory explain state consciousness? The claim of the higher-order theory is that mental

states are conscious when an inner sense produces a second-order representation of them. Being represented by a higher-order mental state is what it is to be a conscious state. Therefore, introspective consciousness explains state consciousness. First-order states represent features of the world, and these are the states that become conscious by means of inner sensing, so state consciousness maintains a mind-world relation on the higher-order inner sense theory. Problem averted.

What continues to bother me about this explanation is that first-order states do not change when they become conscious. The only change that occurs when a state becomes conscious is that a second-order representation appears. And, while a change in relations can be dramatic, the relation of being represented does not seem the appropriate sort of relational change to constitute the difference between unconscious and conscious states. By explaining state consciousness as a matter of being represented by a higher-order mental state, the inner sense theory fails to satisfy the sense that something remarkable has occurred when a mental state becomes conscious.

My solution is to preserve the sensory aspect of inner sensing while eschewing the higher-order structure. Since this higher-order structure has become essential to the notion of 'inner sense', call my theory the *second sense theory*. Think of sensing as a process in two stages, or acts. In the first act of sensing, external senses take various forms of physical stimuli as input and produce sensory states. Then, in a *second* act of sensing, some form of internal sense takes the sensory states as input and produces conscious sensory states. I envision the operation of the second sense as similar to the medieval description of the *sensus communis*:[32] it selects and combines sensory states so as to form a representation of current states of affairs. The second sense theory describes conscious sensory states in terms of a particular sort of relation between conscious states and the world, rather than a particular sort of relation between conscious states and other mental states. In spelling out the particular sort of relation between conscious states and the world, I hope to satisfy the sense that there is indeed a remarkable, yet explicable, difference between unconscious and conscious states.

A central element my account borrows from the inner sense theory is the postulation of a sensory mechanism. One reason this feature is attractive is that it clearly marks the distinction between unconscious sensory states and conscious ones. Conscious sensory states require the operation of a second sense. Without a second sense, no sensory states are conscious. The change from unconscious sensory state to conscious sensory state is dramatic and mysterious. One moment you are busy at the computer, unaware of the cramps developing

in legs and shoulders. The next moment you stop work and are struck by the stiffness and pain. The change from one moment to the next is startling and calls for a comparable change in mental states to satisfy the sense that something remarkable has occurred. Of course, the sense of dramatic change could be mistaken. Simple magic tricks teach us that the seemingly mysterious often submits to mundane explanation. Nonetheless, it would be an advantage if a theory could account for this sense of dramatic change in terms of an identifiable mechanism. A second reason, following from the first, is that the mechanism is empirically testable. This is often cited as a disadvantage of an internal sense because there seems to be no one place where consciousness 'comes together' (Dennett & Kinsbourne 1992). So it seems there could be no mechanism that produces all and only conscious states. As I will argue in Chapter 3, the second sense theory escapes the problematic aspects of Dennett and Kinsbourne's argument. Moreover, I suggest that the sensory mechanism is actually an advantage of the theory. Its empirical element makes it both verifiable and falsifiable. If we can find a mechanism that performs the function attributed to a second sense, this counts in favor of the theory; if not, then we have grounds to reject the theory. In the Appendix I will consider ways in which current empirical research in attention may provide supportive grounds for postulating a second sense.

Third, a second sense theory is sensation-based rather than concept-based.[33] Because conscious sensory states are the result of a sensory rather than cognitive process, they share some important features with sensory states. For example, conscious sensory states may be more detailed than one can conceptually comprehend. A sensation-based theory explains how it is possible to have conscious sensory states about things without knowing anything whatsoever about them, even that they are 'things.'

As a rough characterization, consider three features that seem central to the description of a mechanism as a 'sense': (1) it is non-cognitive; (2) it serves a relay function; and (3) it has the function of producing representations of a particular sort.[34] While these features may not be sufficient, they all do seem to be necessary. I will describe these features here and return to them in Chapter 3 to show how the second sense exhibits all three features.

First, the mechanism in question is *non-cognitive*. As noted, there are important conceptual differences between sensation and cognition, even if there is no clear anatomical division. Foremost, sensation requires no concepts. The minimum requirement for having a concept of a thing is the ability to individuate it, to separate it from its surround.[35] Since it is possible to sense an object without being able to isolate clear borders between that object and its

neighbors, sensation does not meet this requirement. The most I may be able to say about the content of my sensations is some vague gesture, like 'stuff over there.' Nor does sensation require the ability to identify an object or to re-identify it upon repeated presentation, both arguably necessary for cognition. Further, sensory representation relies on tracking, keeping an object in view (or hearing, or touch). Once the object is out of sight, sensory representation ceases. On the other hand, one critical marker of conceptual representation is the ability to represent an object in absence. Thinking about a cup does not require that there be a cup in front of me, but seeing a cup does. In line with this difference, sensory representation is more detailed than conceptual representation. Concepts abstract across differences in particular presentations of an object so as to isolate the key features that will secure re-identification. Because sensations maintain the object in view, abstraction of detail is not required. This brief examination gives us three markers (not to be confused with necessary and/or sufficient conditions) for the distinction between sensation and cognition: sensation is non-conceptual, involves tracking, and is detailed, whereas cognition is conceptual, represents in absence and abstracts from detailed presentations.[36]

A second important feature of a 'sense' is its function as a relay mechanism. External senses take various forms of physical stimuli as input and relay this information in the form of sensory representations to cognitive structures and motor systems. The relay function alone is insufficient to count a process as sensory since cognitive processes also relay information. If we restrict the kind of information processed to non-cognitive forms as described above, then the relay function further specifies the role a sense plays in the mental economy. The job of a sense is to take non-conceptual, tracked, detailed information and pass it on to other systems. We can then further restrict the description of functional role according to specific forms of inputs and outputs. On the input side, the domain of each external sense is specific to a particular form of physical stimulus: the eyes sense light waves, the ears sense sound waves, etc. Each modality has its own form of stimulus to which it alone is sensitive, and therefore each modality is able to represent specific sorts of features. Only the eyes represent colors, and only the touch represents tactile qualities, for example. On the output side, each external sense produces its own variety of representation, it represents the object in a specific way. Eyes produce visual representations, ears produce auditory representations, and so forth.

This description of a 'sense' – a non-cognitive, relay mechanism with specific forms of input and output – will be my basis for taking the second sense to be a 'sense.' While it would be futile to look for a point in neural processing

where sensation gives way to cognition, there are important functional differences between sensation and cognition. My claim in the following is that the second sense deserves to be considered a sense because it shares some of these important functional features with the external senses.

4. Purely verbal?

Suppose I am right that the second sense explains what I have identified (still somewhat broadly) as sensory consciousness. What indication is there that Armstrong and Lycan intend their theory to explain the same phenomenon? Might the dispute here be purely verbal, and higher-order inner sense theorists propose an explanation of what I have called 'self-conscious states'? I am afraid we will have to learn more about the substance of each theory before we can answer this question. At the moment I will begin by offering two indications that the dispute is substantive as well as verbal. (There are always at least verbal disputes where consciousness is concerned, as is already evident in my simple taxonomy.) The first signal of a substantive dispute is Lycan's claim that introspective consciousness explains state consciousness. That is, introspective consciousness constitutes the distinction between conscious mental states and unconscious ones (Lycan 1996: 14). A second, admittedly weaker, indication is the suggestion that introspective consciousness explains 'what it's like' to have mental states (Lycan 1996: 66ff).

As I have said, such terminological choices can have substantive consequences. For instance, in introducing the term 'state/event consciousness', Lycan describes conscious states as states the subject is aware of being in (Lycan 1996: 3). Note that from this description, some form of higher-order theory follows quickly. Here, in shorthand form, is higher-order theorist Rosenthal's argument (Rosenthal 1993b: 356–360, 1991a: 31). A mental state is conscious when a person is conscious of being in it. Having thoughts and having sensations are apparently the only two ways to be conscious of things. So, if a conscious state is a mental state one is conscious of being in, and there are two ways one can be conscious of a mental state, one of those two ways must be the way a mental state is conscious. If a state is conscious by having a thought about it, then the higher-order thought theory is true. If a state is conscious by having a sensation about it, then higher-order perception is true. No other choice seems possible.[37]

Fred Dretske refuses to accept the binary choice offered by higher-order theorists. Instead he disputes the description of a conscious state[38] as a mental

state one is conscious of being in, arguing that its proponents are confusing act and object.

> Either a state is made conscious by (1) its being an *object* or (2) its being an *act* of creature consciousness. A state of creature S is an *object* of creature consciousness by S being conscious of it. A state of S is an *act* of creature consciousness, on the other hand, not by S being aware *of* it, but by S being made aware (so to speak) *with* it – by its occurrence in S making (i.e., constituting) S's awareness of something (e.g., an external object). When state-consciousness is identified with a creature's *acts* of awareness, the creature need not be aware of these states for them to be conscious. What makes them conscious is not S's awareness of them, but their role in making S conscious – typically (in the case of sense perception), of some external object.
>
> (Dretske 1997:6)

Dretske defends the second way of describing state consciousness: A state is conscious when a creature, by virtue of being in that state, is conscious of something in the world. Higher-order theorists describe state consciousness in the first way: A state is conscious when a creature is conscious of being in it.

In Chapter 2 I will argue that Dretske's description of state consciousness in terms of acts is correct. For now I only want to show the consequences that follow from this terminological disagreement. If, in keeping with the higher-order theorist view, conscious states are considered intentional objects, state consciousness is an entirely internal representational affair. Conscious states are those represented by other mental states. State consciousness is not a matter of the way the person is related to the world; it is a matter of the way some of a person's mental states are related to some other of her mental states. On the other hand, if conscious states are viewed as acts through which a person is related to the world, state consciousness is essentially a relation between person and world rather than a relation between mental states.[39] This difference constitutes a substantive disagreement: either mental states are conscious because of the way they represent the world or because of the way they are represented. Therefore, conscious states cannot be considered 'states one is conscious of being in' without begging the question between a higher-order theory and one that is not higher order. Some more neutral characterization is required. I have suggested we start by restricting ourselves to conscious sensory states, which are sensory states that represent the world and that there is something it is like to be in. Higher-order theorists are welcome to suggest an alternate means of identifying conscious sensory states, but they cannot begin with the statement that they are states one is conscious of being in.

One further point is necessary to avoid a purely verbal debate, a point that Dretske overlooks. Sensory consciousness cannot just be a matter of representing the world or there would be no distinction between unconscious sensory states and conscious sensory states. Unconscious sensory states can also represent the world, so more must be said about *how* conscious sensory states represent the world, in what way conscious sensory states are unique, in order to adequately address the question of sensory consciousness. Dretske is careful to distinguish non-epistemic from epistemic forms of sensation, for example, but completely disregards the distinction between unconscious and conscious sensory states. In one blaring example of this disregard Dretske writes: "Seeing a tree, smelling a rose, and feeling a wrinkle is to be (perceptually) aware (conscious of) the tree, the rose, and the wrinkle" (Dretske 1993b: 265).[40] Dretske in no way acknowledges the possibility that these sensory representations could be unconscious in the way that my sensory representation of shoulder and leg cramps are unconscious while I am attending to other matters. Without this distinction, Dretske's 'state consciousness' seems more like Armstrong's 'perceptual consciousness,' the ability to perceive the world (Armstrong & Malcolm 1984: 119). Since perceptual states may or may not be conscious sensory states, a theory of perceptual states cannot explain why some sensory states are conscious and others are not.

In Chapter 3 I will argue that a fuller characterization of the particular sort of representations produced by a second sense can provide such an explanation. But first, in the next chapter I explore the debate between Dretske and the higher-order theorists in order to further clarify the substantive issue lurking beneath the terminological squabbles. Out of this investigation the second sense theory will emerge as a compromise theory, adopting the best aspects of each theory. Following Dretske I argue that sensory consciousness is a matter of representing the world. It is not a matter of sensory states being represented by higher-order states. Following Armstrong and Lycan I recommend a variety of sense mechanism to explain the difference between unconscious and conscious sensory states. Putting these together we have: conscious sensory states are representations of the world produced by a second sense. This description is still incomplete as it does not adequately distinguish conscious world-representations from unconscious ones; nor does it explain what role the second sense plays in their production. To hint of coming attractions, I will account for the distinction between conscious and unconscious sensory states in terms of a representation of the world at the present moment and argue that the second sense is responsible for producing this peculiar form of representation.

Chapter 4 then examines an issue related to, and often conflated with, sensory consciousness, viz. subjectivity. Because conscious sensory states are states of *something*, some description of the subject of sensory consciousness is essential to a full account of the phenomenon. Without such a description, there is a danger that the mystery to be explained by the theory of sensory consciousness simply gets transferred to a mystery about subjectivity. I believe exactly this sort of transference occurs in higher-order thought theories, for example, as I will argue in the next chapter. In such a case, it seems to me we are attempting to explain a phenomenon in terms of a phenomenon that itself requires an explanation. To avoid this consequence, I propose a minimalist theory of subjectivity based on Gibsonian ideas about egocentric space. The form of subjectivity I describe is more fundamental than sensory consciousness, so that a creature could be a subject in this sense without the capacity for sensory consciousness. I contrast this sense of subjectivity with the more sophisticated form of subjectivity described by Lycan, arguing that his theoretically richer account of subjectivity is consistent with – and can be built out of – subjectivity in my minimalist sense.

The final chapter addresses some of the most likely objections to the theory. The fundamental question of whether the theory is too broad or too constrained is addressed in the first two sections. Originally posed by Ned Block as a worry about functionalist theories of the mind, all explanations of mental phenomenon need to face the Scilla of liberalism and the Chrybdis of chauvinism to show that the theory adequately navigates between the opposing threats. Likewise, all theories of consciousness must in some way address what has become known as the Hard Problem: how a phenomenon that is identified non-functionally can be account for in terms of structure and function. Finally, I address objections posed from within the family of representationalist theories of consciousness, one from Rosenthal on the relation between concepts and consciousness and the other from Dennett on the role of reportability in determining what counts as consciousness.

As mentioned earlier, the Appendix aims at making some connections between the neuropsychological literature on attention and the second sense theory of sensory consciousness. I include it as an appendix because it is inessential to the central argument of the book and may not be of particular interest to philosophers. Further, I am not a research psychologist so I am not in a position to draw any definitive conclusions about the connections proposed. This is not to say that these connections are unimportant, however. I firmly believe that our only hope in arriving at a fully satisfying explanation of consciousness is for neuroscientists, philosophers and others to integrate theoretical devel-

opments across domains. Only if we can bridge the gap between mental and physical disciplines do we stand any chance of bridging the gap between mental and physical phenomena.

Chapter 2

On higher-order theories of consciousness

Before we attempt a global synthesis of consciousness theory writ large, to which I adverted at the end of the last chapter, we should begin within the domain of philosophy by considering the reasons for rejecting a higher-order theory of consciousness. According to higher-order theories, conscious states[41] are hierarchically structured such that a higher-order relation between mental states constitutes state consciousness. Two components of higher-order theory are central to the account: (1) one mental state (the higher-order state) takes another mental state (the lower-order state) as its intentional object, and (2) the status of the lower-order state – as an unconscious or conscious state – is determined by its relation to the higher-order state. Advocates of higher-order theory have argued that the second component, the higher-order relation, explains state consciousness. There are two versions of higher-order theory, higher-order thought and higher-order perception theory, and the two differ in some important respects that will appear later. Nonetheless, they share the fundamental claim that a mental state is conscious if and only if a higher-order state takes that state as its intentional object (plus other conditions depending on the variety).

In Part 1 of this chapter I will argue that higher-order theory requires more sophisticated abilities than is needed to explain state consciousness. In higher-order thought theory, the explanandum is a bit elusive, so I will offer three different possible interpretations and worries about each. Some ambiguity also appears in the higher-order perception theory, but here the account seems most plausibly interpreted as a theory of a higher form of consciousness, viz. introspection, rather than an account of the common everyday phenomenon of sensory consciousness. In Part 2 I will present a further challenge to higher-order theories in general based on an objection from Dretske. Though Dretske's objection is not decisive against higher-order theories as it stands, his problem case of Spot does present a puzzle for theories of sensory consciousness.

1. The higher-order explanation of state consciousness

In Chapter 1 I noted that higher-order theorists consider the general category of *state consciousness* to be the mysterious and problematic sort of consciousness requiring explanation. The question as they frame it is: what is it for a mental state to be a conscious state as opposed to an unconscious state? In answering this question higher-order theorists typically begin by describing conscious states as "states one is conscious of being in" (Lycan 1996: 3, 25; Rosenthal 1991a: 31, 1991c: 462, 1993b: 356). A higher-order explanation of state consciousness follows naturally from this description. One is conscious of something when a mental state has an object. Borrowing a grammatical term, we can call this form of consciousness *transitive consciousness.*[42] Transitive consciousness is simply consciousness *of* something. So if conscious states are states we are conscious of being in and consciousness of something is transitive consciousness, then a mental state must be conscious when one is transitively conscious of that state. For a mental state (S) to be conscious, there must be another, higher-order state (HO) that takes S as its object. My mental state S is conscious when I have a higher-order state by virtue of which I am transitively conscious of S. State consciousness is explained in terms of transitive consciousness.

As there are basically two ways to be conscious of something, by thinking about it or by sensing it, there are two corresponding types of higher-order theory. I will elaborate each in turn, first considering the thought-based theory, the *higher-order thought theory*, and then turning to the sensation-based theory, known as the *higher-order perception theory,* or the *inner sense theory.*

1.1 Higher-order thought theory

The clearest description of higher-order structure is given by David Rosenthal in his discussion of the higher-order thought theory. On Rosenthal's account, first-order mental states are sensory, intentional or both. Examples of first-order states are sensations such as feeling a pain or seeing a tree and thoughts such as a thought about lunch or that it is raining. A first-order state (S) is conscious when a higher-order thought takes that state as its intentional object, that is, when I have a thought that I am, myself, in mental state S. A higher-order thought must include a self-referential element in order to identify the particular mental state that the thought is about. "Otherwise," Rosenthal notes, "the thought would just be about that type of mental state, and not the particular token of it. So, in the case at hand, the higher-order thought must be

a thought that one is, oneself, in that mental state" (Rosenthal 1991c:469).[43] The sense of self required need not be very robust, however. All that is needed, Rosenthal says, is "a concept that allows distinguishing between oneself and other things. . .it need not imply that the self has some special sort of unity, or is a center of consciousness, or is transparent to itself, or even that it has mental properties." (Rosenthal 1997:741)

A few more conditions are needed to specify the kind of higher-order thought required. If just any self-referential higher-order thought about a state were sufficient for that state to be conscious, then simply telling blindsight patients about their visual states would be sufficient for those states to be conscious. They would acquire higher-order thoughts, that is, thoughts about their lower-order visual states and – presto! – fully conscious visual states. Rosenthal describes a more familiar case to show the need for further conditions on the nature of higher-order thoughts. If someone tells me I am angry, I may accept this assessment and thereby acquire a thought to the effect that 'I am angry.' But my anger wouldn't necessarily be a conscious state by virtue of this higher-order thought, so we need to be more explicit about what sort of higher-order thoughts are involved. (Rosenthal 1997:737) First, higher-order thoughts must seem *immediate*. Though factors may in fact causally mediate between the lower-order and higher-order thought, no mediating factors can be conscious. The relation between a lower-order state and the thought about it must be seem immediate, even if it is not. As Rosenthal says: "One constraint on the way we are transitively conscious of our conscious mental states is that our transitive consciousness of them *seems*, intuitively, to be immediate. That is, nothing seems to mediate between the mental state itself and our transitive consciousness of it" (Rosenthal 1993b:359. See also Rosenthal 1991a:31; Rosenthal 1997:738).

It is important to note that the higher-order state is not usually a conscious state itself. Only if I have a yet higher-order state (HOT2) that takes the second-order thought (HOT) as its object will the second-order thought (HOT) will be a conscious state.[44] In the same way, the second-order thought (HOT2) becomes conscious if I have a yet higher-order thought. There is a limit to how useful these higher- and higher-order states can be, so there is a limit to how complex the mental hierarchy is likely to become. The limit is empirical and practical, however, not logical.[45]

Finally, higher-order thoughts must be *assertoric* rather than non-assertoric, and *occurrent* rather than dispositional (Rosenthal 1997:742). The higher-order thought must assert – not doubt, wonder or hope – that one is in that state. Further, the higher-order thought must be occurent, according to Rosen-

thal. The higher-order thought must be "roughly contemporaneous" with the state that it is about – not a day before or a week after. "When a mental state is conscious, it is not simply that we are conscious of the state; we are conscious of being in that state. This places a constraint on what the content of these HOTs must be; their content must be that one is, oneself, in that very mental state" (Rosenthal 1997: 741).

Not all higher-order thought theorists agree with the condition that higher-order thoughts be occurrent. In recent years Peter Carruthers has argued for a dispositional theory, where consciousness "consists in a certain sort of intentional content ('analog', or fine-grained), held in a special-purpose short-term memory store in such a way as to be available to higher-order thoughts about the occurrence and nature of those contents" (Carruthers 2000: xiii). States are conscious when available to a system capable of producing higher-order thoughts about their contents. A basic problem with dispositional theories of sensory consciousness was raised in Chapter 1 in connection with Block's description of access consciousness as 'poised.' Sensory consciousness is a property states actually exhibit; it constitutes 'what it's like' to have a mental state. Because sensory consciousness is an occurrent property, a dispositional account is insufficient to explain it. Carruthers recognizes the intuitive force of this objection and offers an intriguing solution. If we adopt some form of consumer semantics (such as the teleo-semantic approach I favor), the content of a representation R is determined, in part, by the use a system makes of that representation. If one of the functions of R is to produce higher-order thoughts, then R acquires a higher-order content whether or not a higher-order thought is in fact produced. For example, the higher-order content a representation of red acquires, according to Carruthers, is "the content *seems red* or *experience of red*" (Carruthers 2000: 242).

Ingenious as it is, this approach fails to account for the difference between unconscious and conscious sensory states. At first glance, the claim that all representations capable of producing a higher-order thought acquire the dual content special to consciousness seems to rule out the possibility of distinguishing between unconscious and conscious sensory states. We need a theory that identifies a property some sensory states have and others lack, and that goes on to offer an explanation of the difference. If all sensory states have higher-order content by virtue of the availability of the state to a higher-order conceptual system, then the presence of higher-order content cannot distinguish between unconscious and conscious sensory states.[46] On closer examination, however, we see that Carruthers divides sensory states into two separate systems, one that is primarily action-guiding and another that is primarily belief-forming

(Carruthers 2000: 152–168). Unconscious sensory states pass through a separate (and presumably, cognitively impenetrable) short-term memory store, and then feed directly into action. These states, therefore, are not capable of producing higher-order thoughts because they form an isolated perceptual system.

I see several problems with this view. First, the crucial work of distinguishing between unconscious and conscious sensations is done by the division into separate perceptual systems with separate short-term stores. I will elaborate on my reservations about this move in a moment, but even granting it, we see that it cannot support a higher-order account. As Carruthers presents the theory, the two systems develop simultaneously as the brain develops, whereas the development of higher-order thought does not occur until much later.[47] Indeed, Carruthers describes the belief-forming system in terms of first-order beliefs about a *rabbit*. So, even if we admit that there are two separate perceptual systems, one to guide action and the other to form beliefs, and that these systems align perfectly with unconscious and conscious sensory states respectively, we would still need further argument that it is the later development of a higher-order conceptual system that renders belief-forming sensory states conscious. But once we have accepted the apparatus of a dual perceptual system, it would seem that a first-order explanation could account for the distinction between conscious and unconscious states in a simpler and more satisfying way.

Moreover, I am disinclined to accept the dual perception theory as described. While Milner and Goodale (1995) propose an intriguing hypothesis that the ventral visual system is responsible for 'off-line' concept forming tasks and the dorsal system is responsible for 'on-line' action guiding tasks, the traditional description of these pathways into object-identification/spatial location (what/where) tasks is still the predominant view. Even if correct, an enormous empirical gap exists between the Milner/Goodale theory of visual processing to the claim that a similar division applies to all perceptual systems and that the division aligns exactly with the division between unconscious and conscious sensory states. I see no evidence for these further claims, and so I see little reason to adopt the dual content higher-order theory on which they rest.[48]

So, putting aside the disposition theory, we now have a description of the higher-order thoughts necessary for state consciousness. Sensory and intentional states are conscious when there is a seemingly immediate, occurrent, assertoric thought that takes that state as its object. A state is conscious when we are conscious of it by having an appropriate thought about it. Stated this technically it is hard to see why the higher-order thought theory would appeal to anyone. How could a mental state be conscious simply in virtue of being the object of a higher-order thought? The answer lies in the representation relation

between the two states. The intuitive idea behind the higher-order thought theory is this: when I think about something, I am conscious of it. When I think about a tree, I am conscious of the tree. When I think about a headache, I am conscious of it too. So perhaps it is by thinking of things that I am conscious of them. My thoughts explain the fact that I am conscious of some things and not others. I am conscious of the things that I am thinking about and not conscious of the things that I am not thinking about. Being conscious of things is *transitive consciousness* and having thoughts is one way to be transitively conscious of things.

The sort of 'things' we are conscious of in the sense of transitive consciousness can be physical or mental. We can be conscious of a physical object like a tree by having a thought about a tree.[49] Alternatively, we can be transitively conscious of mental states like sensations of green or loudness by having thoughts about them. This way of construing transitive consciousness seems to quite similar to, perhaps even identical with, the now familiar notion of mental representation.[50] Consider this passage from Rosenthal:

> Being transitively conscious of something means that one is in a mental state that represents that thing. Moreover, something close to the converse holds as well. It is obvious that when we have thoughts about things, we are transitively conscious of those things. And sensory states, such as pains and perceptual sensations, also typically represent qualities and properties of one's physical environment or one's own body. All mental states, moreover, are either intentional or sensory. So being in a mental state is very often sufficient for one to be transitively conscious of something. (Rosenthal 1993b: 356)

Rosenthal resists saying that transitive consciousness is nothing more than mental representation by limiting the converse entailment from mental representation to transitive consciousness. If being in a mental state is 'very often' sufficient for transitive consciousness, and the key feature is their 'aboutness' or their representational nature, then why hold back?[51] Why not simply identify transitive consciousness with mental representation? Rosenthal seems to want it both ways. He wants to invoke the relative theoretical security of mental representation to explain the mysterious and ineffable phenomenon of state consciousness. Yet he wants to retain the notion that transitive consciousness is a kind of *consciousness*, not just some plain old mental representation. But by resisting the identification of transitive consciousness and mental representation, Rosenthal leaves himself open to the objection that the notion of transitive consciousness retains a residual air of mystery not fully explained by the higher-order thought theory .

Here is the problem. If transitive consciousness is merely mental representation, it is no longer obvious how we are conscious of the things we think about. Let me dwell on this point for a moment because it applies to both sorts of higher-order theory.[52] The higher-order thought theory is built on the intuitive notion that we are conscious of things when we think of them. However, the force of this intuition seems to rely on the idea that the thoughts by which we are conscious of things are themselves conscious states. If I am in an unconscious state, such as dreamless sleep, it is not at all obvious to me that the occurrence of a mental representation of a tree will make me conscious of that tree in any interesting sense at all.

Rosenthal focuses his argument on the explanation of state consciousness and does not elaborate on how transitive consciousness diverges (if it does) from mental representation. With this sort of sleight of hand higher-order theorists are able to import the sense of state consciousness to be explained into the notion of transitive consciousness. The phrase 'conscious of x' rather than the more simple 'represents x' subtly lends plausibility to an otherwise implausible description. The question to ask oneself when considering a higher-order theory of state consciousness is: where is the mysterious aspect of 'consciousness' located and how has it been explained? To begin, the mystery is situated squarely in *state consciousness*. But I suspect it will shift before the story is over.

In an effort to avoid some of the confusion generated by the notion of transitive consciousness, I will not use the term at all except in the context of describing higher-order theory. In this context I will be careful to use the technical 'transitive conscious of' to mark this peculiar form of representation. Further, I will distinguish the case where one is 'transitively conscious of' something from my own technical sense of being 'state conscious of' something. In my terminology, one is state conscious of something just in case one has a *conscious state* that represents that thing. A person who has only unconscious mental representations is not *state conscious of* anything, although she may be *transitively conscious of* all sorts of things.

To return to the higher-order thought theory, the claim is that, just as we are transitively conscious of things in the world when we think about them, we are transitively conscious of our mental states in the same way. By having a thought about a mental state we are transitively conscious of that mental state, and being transitively conscious of a mental state is what it is to have a conscious mental state. On the higher-order thought account, state consciousness is a matter of transitive consciousness of that mental state.

1.2 Objections to the higher-order thought theory or playing the shell game

In formulating my concerns about the higher-order thought theory, I had the frustrating sense that the phenomenon to be explained kept shifting under my gaze. First the mystery was in conscious states, then it seemed to shift to the subject of consciousness and later to the higher-order thought itself. So my overall objection to the theory is similar to equivocation, except that the ambiguity is not in the meaning of the words but the location of the explanandum. Call this the *Shell Game Argument*.

At first, as described above, the mystery of 'consciousness' is located squarely in the lower-order state. The first question, then, is how a state could be conscious merely in virtue of bearing a representation relation, such as being thought about. One way to put the problem is: why are mental states conscious when thought about, but other things, like tables, are not?[53] Rosenthal answers this objection analytically. Higher-order thoughts do not *cause* a mental state to be conscious; rather, "a mental state's being intransitively conscious[54] simply consists in one's being transitively conscious of it" (Rosenthal 1997:739). Lycan gives a similar response in defending the higher-order perception theory from the same objection. Mental states, but not tables, are the sort of thing that we can be transitively conscious of being in. We use the word 'conscious' to mark the distinction between the mental states we are transitively conscious of and those we are not transitively conscious of at any given moment (Lycan 1996:24). Saying that only mental states can be conscious is like saying that only women can be sisters. No matter how the sibling relation is applied to men, they can never be sisters (except perhaps metaphorically). We only apply the word 'sister' to women. Similarly, a table cannot be conscious when thought about because state consciousness only applies to mental states.

But this response raises a second problem. The acquisition of a female sibling does not seem appropriately analogous to the acquisition of a conscious state. Over time the acquisition of a new sister will have many effects, but the relation itself makes no significant or mysterious difference in the nature of the one who acquires it. The difference between unconscious and conscious sensory states *is* significant and mysterious, however. In one case there is nothing it is like to have a sensory state and in the other case there is something it is like to have that state. The relation of being represented does not seem to be sufficient to constitute this sort of difference. My therapist may represent my mental states in all sorts of ways without any remarkable difference in the states so represented. I may even represent my own states with no remarkable change, as

in Rosenthal's example of acquiring a thought about my anger. The difference between unconscious and conscious sensory state is remarkable, but acquiring the relational property of being represented does not seem to constitute the right sort of difference.

This is not to say that state consciousness cannot be a relational property of some sort. We sometimes discover that seemingly intrinsic properties, such as weight, are in fact relational. Given sufficient theoretical background information, we can understand how force and mass combine to yield weight and how a difference in force yields a difference in the weight of an object in sometimes surprising ways. Still, it is difficult to see what sort of background theory could help us understand how the relation of being represented could constitute the difference between an unconscious state and a conscious state. In no other case does the relation of being represented constitute such a striking difference.

Higher-order theorists accept this consequence and even consider it a benefit of higher-order theory. First-order states remain the same state, whether they are unconscious or conscious. Unconscious pain is still pain, even though unconscious, as is evidenced by the important causal relations it has in common with conscious pain, such as wincing or avoidance behavior. As examples of unconscious states, Rosenthal cites cases of mental states that are conscious intermittently, like a headache that comes and goes, or the 'light bulb' phenomenon where the solution to a puzzle suddenly pops into your mind without any conscious deliberation (Rosenthal 1997:731). What explains these cases, Rosenthal argues, is that unconscious states are conscious when there is a higher-order thought about them. The same headache remains throughout; it is conscious intermittently by virtue of periodic higher-order thoughts. Similarly, the puzzle continues to be the object of unconscious thoughts until finally the solution becomes conscious by virtue of a higher-order thought about it.[55] Mental states remain the same states when conscious, according to higher-order theorists; they are not changed at all by their higher-order representation.[56]

So, what accounts for the remarkable change when a state becomes conscious? Look under Shell 2. For the higher-order thought theory, the change is in the creature, not in the mental state. By thinking (in the right way) about mental state S, I become transitively conscious of S. What it is for S to become conscious is not that anything startling happens to *S* – nothing happens to S – but now *I* have a thought (of the right kind) about it. I am now transitively conscious of S. This interpretation is supported by Rosenthal's restriction that transitive consciousness applies only to creatures, not mental states.[57] He writes:

> Being transitively conscious of something is a relation that a person or other creature bears to that thing. So only creatures can be transitively conscious of things. A mental state may well be that state in virtue of which somebody is conscious of a thing, but the state cannot itself literally be conscious of anything. (Rosenthal 1997:738)

Even though a mental state is conscious by virtue of a higher-order thought about it, the critical element is that *I* am the one who has the appropriate thought. Where, then, has the mysterious aspect of conscious states gone? It has shifted to me. But who am I and what is it for me to be conscious? It won't do to say, as Rosenthal does regarding the self-reference of conscious states, that only a minimal concept of self is required. While the ability to "distinguish between oneself and other things" may be sufficient to get one's higher-order thoughts to refer to oneself (Rosenthal 1997:741), self-reference alone does not specify what sort of creature is capable of being transitively conscious of things or what sort of change occurs in the creature. According to Rosenthal, creature consciousness is not especially mysterious, but it seems that in the end the mystery of state consciousness lands squarely on the creature, albeit in the form of transitive consciousness. If this interpretation is correct, Rosenthal needs to develop an account of the self that is explanatorily more basic than state consciousness. I offer the beginnings of such an account in Chapter 4, but the form of subjectivity I discuss will be too basic to accommodate the higher-order thought theory. Whether a sufficiently robust sense of self could be developed out of this account remains to be seen.

An alternative interpretation is that the mysterious aspect of conscious states has shifted to the higher-order thought itself: Shell 3. In this version of the theory, the whole higher-order thought is conscious. Rocco Gennaro has defended this description of the higher-order thought theory, arguing that conscious states are *complex* states composed of a lower-order state and its higher-order representation (Gennaro 1996:16).[58] Gennaro calls this the Wide Intrinsicality View (WIV): conscious states are individuated widely to include both lower-order state and higher-order thought, and state consciousness is intrinsic to the resulting complex. This move has the immediate disadvantage of obscuring the neat hierarchical arrangements of Rosenthal's theory. On Rosenthal's view the lower-order state is conscious – there is something it is like to be in it – when it is represented by a higher-order thought. But on Gennaro's theory, it is unclear whether the whole complex state is conscious and there is something it is like to be in only part of it, or whether there something it is like to be in the whole complex state. If the mystery resides in the whole state, then

the theory is circular. The original description of higher-order thought theory avoids circularity by explaining one mysterious form of consciousness – state consciousness – in terms of another form – transitive consciousness. But if the mystery now resides in the complex state of higher-order thought plus lower-order state, the higher-order thought is explaining itself. On the other hand, if the mystery resides in part, the situation is no better. Either the lower-order state is the conscious part and the mystery has simply shifted back to the position where it began, or the higher-order thought is the conscious part and the theory again is circular. In any case, the introduction of such a view confirms my suspicion that the mysterious aspect of conscious states has shifted disturbingly from lower-order state to higher-order thought or the creature who bears it.

To summarize the moves in this funhouse Shell Game, we begin with the mystery in the lower-order state. In this case, a more substantive account of *transitive consciousness* is needed. How does transitive consciousness differ from mental representation? And given lower-order states do not change when represented, how does transitive consciousness make it like something for the creature to have those states? If the mystery is in the creature (Shell 2), a substantive account of the creature is needed, such that it is the kind of creature capable of conscious states *and* that its having higher-order thoughts is sufficient to constitute the sort of dramatic change that occurs when states are conscious. If the mystery is in the higher-order thought (Shell 3), then the theory needs to explain whether it is like something to be in the whole higher-order state or only part. If the whole, then the theory is circular. The higher-order state is explaining itself. If part, then either the lower-order state is what is conscious and we are back to the original position, or the higher-order thought part is conscious and the theory again is circular.

Assuming higher-order thought theorists can clear up this difficulty, they still face another problem. If we are to explain state consciousness in terms of thoughts, we need some description of what thoughts are. Rosenthal tends to be slippery in discussing thoughts, and for good reason given the sort of work he needs them to do. On one hand thoughts need to have a fairly well-developed cognitive structure in order to represent both a mental state *and* that the bearer of that mental state is at that very moment having it. On the other hand, infants and non-human animals arguably have conscious states without benefit of such well-developed cognitive abilities. To accommodate creatures on the low end of the intellectual scale, Rosenthal minimizes the requirements for thoughts.

> To have thoughts a command of some concepts is of course necessary. But those conceptual abilities need not be anything like as powerful as ours, nor need the concepts be nearly so fine grained. (Rosenthal 1997:741)

> [O]ne need not have much ability to think to be able to have a thought that one is in a particular sensation. Infants and non-human animals can discriminate among external objects, and master regularities pertaining to them. So most of these beings can presumably form thoughts about such objects . . . No more is needed to have thoughts about one's more salient sensory experiences. (Rosenthal 1991c:472)

When talking about sensory states, Rosenthal's description of thoughts gets so loose that it is sufficient to *point* to a sensation in the visual field in order to refer to that sensation in thought.

> We refer in thought to physical objects by way of their position in our visual field. It is natural to suppose that a thought can similarly refer to sensory states by way of their position in the relevant sensory field. In any case, conscious differentiation of sensory detail quickly outstrips one's conceptual resources; so some such means of referring to sensory states is necessary. (Rosenthal 1997:741f)

But this statement goes too far. To say that conscious differentiation outstrips conceptual resources undermines the principal advantage of the higher-order thought theory.[59] Higher-order thoughts are likely necessary to conscious states, Rosenthal argues, because a difference in thought content changes the content of conscious sensory states. When wine tasters and music lovers acquire new concepts such as 'tannic' and 'atonal' through training, these concepts allow finer-grained awareness of sensory qualities (Rosenthal 1991c:472, 1991a:34, 1997:742). Acquiring concepts allows connoisseurs to differentiate more of the qualities available in their sensory states. Before they acquired these concepts, their sensory states were just the same states as afterward, but without the concepts, the sensory qualities contained in those states remained unconscious. Seems reasonable, but I wonder, how is it that the information in some sensory states admits of pointing but the information in others does not? I can point to that new shade of blue in the upper left quadrant of my visual field and so be conscious of that peculiar blue sensation but I cannot point to tannic or atonal elements in these sensory states so as to be conscious of them? Why not?

More troubling, how does one point to a sensory state one is not conscious of being in? On the higher-order thought view, one is pointing to a feature of one's *sensory state*, not a feature of the world. Not only does this description assume a rich, detailed visual field projected somewhere inside the

mind, at which one can point,[60] this visual field must be *unconscious*. On the higher-order thought theory, features of a sensory state are conscious when represented by a higher-order thought. In the passage above Rosenthal suggests that one acquires thoughts about novel sensations by means of pointing to them. The act of pointing to the sensation secures demonstrative reference, and thereby a thought is formed. By virtue of the thought, the sensation is conscious. Thus the pointing must be *prior* to state consciousness, if only momentarily. Thus one is pointing somehow to a sensation in an internal visual field one is not conscious of being in. Very odd.

The whole idea of pointing to sensations in the sensory field seems ad hoc, a wily attempt to avoid the Myth of the Given (Sellars 1963). This is the problem Rosenthal faces. We sometimes have conscious sensory states that we have never had before, such as seeing a new shade of the sky or encountering indigenous spices or unfamiliar noises in a foreign country. Since concepts require some form of learning, some process of acquisition, we have no concepts of such novel sensations. But on the higher-order thought account we need a concept sufficient to refer to a sensory state in order to for it to be conscious. Unless there is some process of concept acquisition that applies to these cases, either we cannot have conscious sensory states about novel things, which we plainly do have, or higher-order thoughts become a form of the Given. Non-inferential knowledge of the content of our sensory states would be given in the act of sensing. For my novel sensation of blue to be a conscious state I would need to have a thought to the effect that 'I am now in a visual state of new blue.' Without some other way of referring to visual states, facts about them, such as the fact that I am now in visual state X (new blue), would need to be given in the act of sensing in order to have the higher-order thought by which the sensation becomes conscious.

To avoid this conclusion, which no doubt Rosenthal would like to do, he needs to find a way to refer to novel sensory states in the absence of a concept specific to them. The ingenious solution is demonstrative reference. If we can point to a sensation and identify it at least rudimentarily as *that*, we can thereby have a higher-order thought to the effect that 'I am now in *that* mental state', where 'that' refers to the sensation of new blue. Referring to sensory states by pointing would solve another problem for higher-order thought theory as well. As noted above, non-linguistic creatures seem to have conscious sensory states despite their rudimentary cognitive abilities, and indeed we seem able to have conscious sensory states for which we have no words, like new blue. As Rosenthal says, our conscious sensory discriminations outstrip our conceptual resources. So maybe we can just point to our sensory states.

But again I wonder, how do you point to a sensory state? When referring to an external object in thought, you point to it by using some sort of sensory tracking (Clark 2000: Chapter 4; Millikan 1993: 271; Evans 1982: 150). You listen to a peculiar creaking in the back of the house or follow a lovely butterfly with your eyes. Sensory tracking cannot by used to identify a sensory state, however, since per the higher-order thought hypothesis, state consciousness involves no inner sense. By what means, then, is pointing accomplished? No theory is offered. The suggestion of an inner sense raises the further question of whether demonstrative reference is best accomplished by a 'thought'. Thoughts require concepts, even if they are minimal ones of sensory demonstratives. Why do we need even this sophisticated a cognitive act for simple demonstrative reference? Rather, suppose that internal demonstrative reference is secured in just the way external demonstrative reference is secured. Suppose there is an inner sense by which we pick out our sensory states. The proposal of an inner sense, by adding a mechanism for internal sensory reference, has the virtue of reducing the cognitive requirement for higher-order states. Higher-order states might be as concept-free as sensory states. The inner sense theory allows us to be transitively conscious of sensations we have no concepts about and allows creatures with limited (or no) conceptual ability to have conscious states of any kind. The addition of a sensory mechanism to secure demonstrative reference seems a decided advantage in favor of the inner sense theory, but in the next section we will see that higher-order inner sensing faces problems of its own.

1.3 Higher-order perception theory[61]

The basic representational structure of higher-order thought theory also applies to higher-order perception theory. A lower-order state is conscious when represented by a higher-order state. In the higher-order perception theory, the higher-order state is formed when an inner sense scans the lower-order state just as an external sense might scan an object. Lower-order states are *sensed by* rather than *thought about by* the higher-order states. Still, the intuitions that motivate higher-order perception theories are the same as for higher-order thought theories. We become transitively conscious of things in the world by sensing them using our external senses, so it is reasonable to suppose we become transitively conscious of internal things like mental states by sensing them with an internal sense.

The structure of the higher-order perception theory is clearest in Armstrong's description of the levels of consciousness. The level comparable to Rosenthal's first-order states is *perceptual consciousness*, the ability to perceive

the physical world. At this level the person may have no awareness at all of his perceptions, as Armstrong puts it, but is nonetheless able to respond appropriately to perceptual information (Armstrong & Malcolm 1984: 119).[62] Armstrong's most famous example of perceptual consciousness is the truck driver who suddenly realizes he was unaware of driving the last 10 miles (Armstrong 1981: 59). Though able to respond to sensory stimuli by making turns, accelerating up a hill, and the like, the driver is not aware of these responses or the perceptions that make them possible. Sleepwalkers are another example:

> Although they are perceiving, and although, clearly, they have purposes of sorts, if only the purpose to walk down the stairs, they appear not to be aware that they are perceiving and appear not to be aware of what it is that they are aiming at. They appear to lack *introspective consciousness*.
>
> (Armstrong & Malcolm 1984: 120)

To be aware of our perceptions requires the higher level of introspective consciousness. As Lycan puts it, introspective consciousness is "'perception' by Lockean 'inner sense,' i.e., by focusing one's attention on the internal character of one's experience itself" (Lycan 1996: 4). As with higher-order thought theory, conscious states are described by higher-order perception theorists as mental states one is conscious of being in (Lycan 1996: 25). A mental state is conscious when scanned by some sort of inner sense which produces a higher-order sensory representation of that state. To use the terminology introduced earlier, introspective consciousness is a variety of transitive consciousness that explains state consciousness (Lycan 1996: 4). The basic representational structure common to both higher-order theories is this: the higher-order state represents the lower-order state, and by virtue of this higher-order representation, the lower-order state is conscious.

Despite this common structure, some key terminology differs between the two higher-order theories. Recall that for higher-order thought theorists, 'introspection' is when a higher-order thought is conscious by virtue of a third-order thought. According to Rosenthal, higher-order thoughts are simply one form of *transitive consciousness*, not a case of introspection at all (Rosenthal 1997: 745). For higher-order perception theorists, on the other hand, *all* higher-order representation produced by inner sensing is *introspective consciousness*. There can be layers of higher-order representations, as well as either passive or deliberate acts of introspection, but all are forms of introspective consciousness. Since higher-order theorists of both types agree on the representational structure of state consciousness, this terminological discrepancy may seem unimportant. But the characterization of introspection as a third-

order representation by higher-order thought theory, on the one hand, and as a second-order representation by higher-order perception theory, on the other hand, points to a general confusion about the distinction between state consciousness and introspection.[63] Before we can answer the question of state consciousness, we need to come to some agreement about these distinctions or, as seems to happen to higher-order perception theorists, state consciousness will tend to be conflated with introspection.

1.4 Objections to the higher-order perception theory

On either higher-order theory a mental state becomes conscious without undergoing any fundamentally internal change. The mental state acquires a representation relation – it is now represented by another mental state – but otherwise it remains the same. So as with higher-order thought theory, we can ask how a mental state could be conscious without undergoing some sort of change. The higher-order perception theory opts for Shell 2: the change occurs in the bearer of the mental state, not in the state itself. By sensing mental state S with my inner sensor, I become conscious of it. The change is in me, not in the state. Higher-order perception theorist William Lycan explicitly places the mysterious aspect of consciousness at this higher representational level, called 'subjective consciousness':

> This is (metaphorically speaking) having a "point of view." A subject's consciousness in this sense is "what it is like" for the subject to be in whatever mental states it is in. A tighter characterization might be: what can be described, if at all, only in the first person. (Lycan 1996:4)

Now the trick is to figure out who the subject is and how she comes to have a 'point of view'.

A problematic way to interpret the higher-order perception theory is to suppose that external objects are presented by the senses in a sort of inner theater and an 'inner eye' views the objects as presented on an inner screen. On such a view a mental state is conscious when it appears on screen and is viewed by the inner eye. But since there is in fact no inner theater and nothing analogous to an inner eye, the account is empirically false. The problem of course is in the very first step of supposing the senses present inner objects for display. If there were such inner objects, we would need to know what the relation is between the inner and external objects. Perhaps the relation is resemblance: the inner objects have similar features to external objects. Then we would need to know which features are similar. Certainly size and weight could not be re-

produced inside the head. Perhaps shape and color could be replicated in some way. But even here we face the empirical problem that nothing inside the head is shaped or colored in the way that external objects are shaped and colored. So the idea of 'inner objects' requires a more subtle account if it is to be consistent with what we know about the brain.

A far more plausible account is that sensory states represent external objects in some way other than recreating the object in miniature within the head. How does sensory representation work? How do sensory states represent objects? The answer depends on your theory of mental representation generally, which is a matter of debate in itself. The higher-order perception theorist need not defend any theory in particular. In whatever way mental states manage to represent objects, their claim is that sensory states represent features of objects in that way. Higher-order states in turn represent sensory states in the same sort of way.[64]

If state consciousness is a matter of higher-order representation of sensory states, a question arises that proves particularly illuminating in regard to higher-order perception theory: what is the function of higher-order states? First order states are conscious when represented by higher-order states. But why bother with higher-order representation when unrepresented (and so unconscious) first-order states are doing the critical work of representing the world and its features? Dretske states this objection bluntly: "HO [higher-order] theories of state-consciousness make questions about the function of consciousness hard to answer. Or, worse, they make the answer obvious: it has no function" (Dretske 1995: 117).

The situation may not be so dire as all this, however. Representing external objects and their features is not the only option; there are several possible internal functions that higher-order representations might perform. Higher-order thought theorist Gennaro suggests that conscious states might produce beliefs or behavior that unconscious states would not (Gennaro 1996: 10). This suggestion seems plausible until you ask what beliefs and behavior could not be included among the causal relations of first-order states. Since first-order states need not be conscious to affect cognitive processes (as shown in the blindsight cases and covert processing studies described in Chapter 1), there seems no reason that first-order states could not directly cause all the beliefs and behavior that involve their content. An exception might be the belief that one is in a conscious state or behavior like saying "I am conscious". Yet even these effects could possibly be produced without there being any actual conscious states. One of the problems in determining whether robots have conscious states is the question of whether and when to trust the robot's reports. A robot could

be programmed to say 'I am conscious,' and maybe even to believe it, without really being in a conscious state. While there may be ways to determine when a report like 'I am conscious' is genuine, it is unlikely that the *function* of state consciousness is to produce such reports.[65]

Higher-order perception theorists have a more promising proposal. Lycan suggests that higher-order representations serve to coordinate and relay lower-order representations (Lycan 1996: 14).[66] Armstrong originally put forward the idea that inner sensing integrates mental states to get them working together in complex ways.

> The co-ordination can only be achieved if the portion of the computing space made available for administering the overall plan is continuously made "aware" of the current mental state of play with respect to the lower-level operations that are running in parallel. Only with this feedback is control possible.
> (Armstrong 1981: 65)

It is now fairly well established that much of the processing of sensory stimuli occurs through the operation of parallel, distributed systems. Inner sensors might monitor these sub-processors, producing representations of their states so as to control action. In order to, as Armstrong says, 'administer the overall plan', an organism must be aware of what is happening to it at the moment. In order to reach for the red glass in front of me and take a drink of water, I need to be kept abreast of the location of my hand relative to the red glass, the feel of the glass as my fingers grip its slick sides, and the shifting weight of the water as it rolls toward my mouth. All of this sensory information needs to be coordinated in order to complete the intended action.

But wait. The information I just described is not information about my *sensory states*, it is information about things in the world: my hand, the glass, its texture, the water. The sort of coordination I just described is coordination of information about the world, because knowing about the way the world is, in the first instance, what I need to know to act appropriately in it. So how might information about my *sensory states* aid in behavioral control? I see two possible interpretations of this claim, which I will develop below. Though I think the second interpretation is a better description of the Lycan/Armstrong theory, the first is a live option and is sometimes taken to be the claim of the higher-order perception theory.[67]

The first possibility is to say that information about the world is made available *by virtue of* information about sensory states. By representing my sensations of hand motion, the red look and slick feel of the glass, etc. I thereby represent my hand, the glass, etc. It is certainly the case that by virtue of sensory

states themselves I represent external objects, for this is just to say that sensation is one form of mental representation. But this is not the same as saying that by virtue of *representing my sensory states* I represent external objects, and this is what is under consideration as a possible interpretation of the higher-order perception theory. This thesis is tantamount to claiming that the content of a representation, what a sensation represents, can be read off the vehicle of representation, the sensory states. By sensing my sensory states I come to represent their content.

On the face of it, the proposal sounds like good-old sense data theory, and the charge of invoking sensa has been leveled against higher-order perception theory before (Güzeldere 1995:353). Posited to account for the way physical objects appear to a perceiver, sensa are objects that have the properties physical objects appear to have (Broad 1927/1969:240). So if it appears to me there is a red glass in front of me, then there genuinely is a smooth, red, cylinder-shaped sensum by virtue of which I perceive the glass. One problem with this view is that physical objects have determinate properties – it either is or is not the case that this glass is cylindrical. But it seems that sensa could be indeterminate. For example, if I saw the glass from a distance, I might not be sure whether it is cylindrical or not; the glass appears neither cylindrical nor non-cylindrical. So, given that sensa have the properties they appear to have, my glass sensum would be neither cylindrical nor non-cylindrical. Nor could the issue be resolved. If I move closer and determine that actually the glass is cylindrical, then either the sensum has more properties than it appeared to have or this new inference is based on an entirely new set of sensa that includes determinate cylindricality. If the former, then the theory has not bridged the appearance/reality gap it was proposed to bridge. If the latter, then sensa have the strange feature of being objects about which nothing additional can be learned.[68] Sensa would be very strange objects indeed.

As already noted, contemporary theories of representation eliminate the need to posit metaphysical objects such as sensa to serve as intermediaries for representing physical things. Lycan, for one, has clearly rejected the notion of sensa, stating explicitly that the objects of higher-order states are intentional objects, not metaphysical ones (Lycan 1996:71). If sensory states are representations of external objects or their features, then some sort of relation must accomplish the representation. But the relation need not be one of literal resemblance or picturing; sensations need not look like what they represent, nor need they have the properties they represent objects to have. Higher-order perception theorists only need to claim that there exists an appro-

priate isomorphism between sensory states and the objects (or features) they represent.[69]

The critical question is how the inner sensor might get a hold of that isomorphism if it is to capture the content of the sensory representation in the higher-order state. In order for the higher-order state to represent the world and its features *by virtue of representing* sensory states, the inner sense must be able to determine what the sensory state represents. Given the teleo-functional bent of Lycan's work, he would likely answer this question in terms of function. Sensory states and scanners could be designed so that scanning the sensory state has the function of producing representations about the world. By scanning my sensations of red and slick, I thereby represent the glass. By scanning my bodily sensations of hand motion and combining these with my visual sensations of peach-colored hand-expanse, I thereby represent my hand moving toward the glass. Integrating all of these together I can reach for the glass and take a drink of water.

I see two problems with this scenario. First, it is not higher-order. The integrated representations described are about the world – the glass, my hand – not the sensory state. The sensation drops out of the final representation, as any good vehicle should. Stressing the importance of the vehicle/content distinction, Dretske provides the useful analogy of words (vehicle) and story (content). The words are in the book, but the story, the dragons and dwarfs and fairies, are not in the book (Dretske 1995:34f). In order to tell a story about dragons, dwarfs and fairies, you have to use words. The words are necessary, but it is not by *representing* the words that you represent the story. Similarly, sensations are necessary to representing objects, but it is not by *representing* sensations that you represent objects. This point raises the second problem: higher-order representation is unnecessary in this scenario. There is simply no need to represent sensory states in order to represent their content. The content is there already in the sensory state, so why duplicate it at a higher level? If we need to integrate several representations in order to coordinate action, we can simply *conjoin* them. There is no need for an extra level of representation to do this job. So, even if we can make sense of how higher-order states represent sensory states, we need a better explanation of why higher-order representation is necessary to sensory integration and behavioral control. [70]

Another possibility open to the higher-order perception theorist is suggested by the term 'introspective consciousness.' Perhaps higher-order states represent sensory states *as* sensory states rather than as content-bearers. Instead of representing my sensory states in order to represent the world, I represent my sensory states to find out what is happening with my sensory states. I focus

on the way I am sensing rather than on what the sensation represents. So in scanning my sensory states while reaching for the glass, I represent my sensation of red and slick in order to keep track of what is going on inside of me, not to find out anything whatsoever about the world. Certainly information about how I am sensing can be useful in many contexts, and the process of internal assessment involved in gaining this sort of information is suitably called 'introspective consciousness'. Moreover, this seems to be what Armstrong and Lycan have in mind. Armstrong speaks of introspective consciousness as "not directed toward our current environment and/or our current bodily state. It is perception of the mental" (Armstrong 1981:60). Introspective consciousness is a matter of "representing that certain perceptual and other representations are currently part of our mental life." (Armstrong, TW: 5) Finally, introspective consciousness allows discrimination among thoughts and feelings (Armstrong 1981:14). Lycan for his part speaks of introspective consciousness as "deliberate scrutiny of the contents of one's phenomenal field . . . To be actively-introspectively aware that P is for one to have an internal scanner in working order that is operating on some state that is itself psychological and delivering information about that state to one's executive control unit" (Lycan 1987:72).

My problem with this description of higher-order perception theory is that it is too sophisticated an activity to be simple sensory consciousness. As social creatures, humans have indeed come to represent our sensory states in order to talk about them. We have great fun contemplating the possibility of strange sensory arrangements such as inverted spectra. Talk about sensory states is integral to the analysis of aesthetic judgements, and competent self-evaluation of our visual states is necessary to proper optometric care (blurry or clear?). But all these examples are, well, higher-order. They hardly seem consistent with the sort of run of the mill conscious sensory states of my cat, for instance. It is perfectly intelligible to imagine creatures capable of wakeful, alert conscious sensory states that, sad to say, never achieve the ability to represent their sensory states *as* sensory states. Indeed, my cat, along with most or all non-human animals, probably does not have this ability. Likewise, developmental studies suggest that very young children are incapable of higher-order states as well. Children under the age of three have difficulty distinguishing a real object from a thought about an object (Wellman 1988), and not until about four years of age does the child understand that her mental representations about things may be fallible or may differ from the representations of others (Tomassello 1993).

All of this indicates that the sensory consciousness of animals and infants (as well as adults most of the time) is a matter of representing the world, not

sensory states. Why and how conscious sensory states represent the world and what distinguishes them from unconscious sensory states are all good questions, and I will do my best at answering them in the next chapter. But we should be clear about the difference between sensory consciousness and introspection. Sensory consciousness is an everyday, run-of-the-mill occurrence. It is marvelous and miraculous if you think about it, but most often we don't. We wake up in the morning and move around in the world, and only in unusual circumstances do we introspect the sensory states by which we navigate. What is marvelous and miraculous is the way conscious sensory states represent the world, not the way in which they are represented by other mental states. Though adult humans do have the ability to represent our sensory states, we exercise that ability rarely. Sensory consciousness, however, occurs every waking hour. Therefore introspective consciousness cannot explain sensory consciousness.

Rather than explaining sensory consciousness, introspection is a different form of consciousness altogether. When you follow Lycan's instructions and "concentrate on your visual field as such, then focus on a particular patch of red, then shift your attention to the upper left quadrant of the field," you are no doubt in a mental state that represents one of your sensory states, viz. the sensation of red in the upper left (Lycan 1996: 17). On a normal day, however, when you are not reading philosophy, you are far more likely to concentrate on the field in the world, and focus on a particular patch of red tomatoes. If so, you are in a conscious sensory state that represents the world. And these conscious sensory states occur regularly without any act of introspection.

2. Formulating an alternative: A flat theory of sensory consciousness

According to the argument of the first section, a cognitive theory such as the higher-order thought approach is problematic because it is unclear how to secure demonstrative reference to sensory states without some sort of sense mechanism. But the theory that provides such a mechanism, the higher-order inner sense theory, seems to be a theory about introspection rather than a theory about sensory consciousness. So perhaps the solution is to consider a new sort of sense, call it the *second sense*, that produces representations of objects in the world. The second sense takes sensory states as input and produces conscious sensory states as output. We might call this a '*flat*' rather than 'higher-order' theory because sensory consciousness is explained in terms of a mind-world relation rather than a mind-mind relation.[71] On a flat theory, a mental

state is conscious when the creature, by virtue of being in that state, has a particular sort of representation about the world. On a higher-order theory, by contrast, a mental state is conscious when a creature has a particular sort of representation about *it*.

Some (notably Carruthers 2000) have used the term 'first-order' theory to denote non-higher-order accounts of consciousness. At the risk of swimming against the tide, I have chosen to use my own term for two reasons. First, states other than first-order states can be conscious, and the flat account ought to be able to extend to them. At least, I see no reason to terminologically restrict possible extensions of the theory. Second, my view differs significantly from first-order accounts given by Dretske and Tye in attempting to offer an account of the difference between conscious and unconscious sensory states. Dretske, as I argue later in the chapter, does not clearly address this difference at all, so a fair amount of textual interpretation is required to determine his view. Tye, as I argued in Chapter 1, offers a cogent answer to the question of qualitative character – the difference between sensory representations of red and sensory representations of green – but also does not clearly address the difference between conscious and unconscious sensory states. Nonetheless, my view continues in the line of reasoning begun by these theorists in arguing that sensory consciousness should not be explained in terms of higher-order relations.

In the first part of the chapter we considered some of the problems in higher-order accounts, but we need an alternative account if we are to reject the higher-order view. Fred Dretske, the most vocal opponent of higher-order theory, presents the alternative in this way:

> Either a state is made conscious by (1) its being an *object* or (2) its being an *act* of creature consciousness. A state of creature S is an *object* of creature consciousness by S being conscious of it. A state of S is an *act* of creature consciousness, on the other hand, not by S being aware *of* it, but by S being made aware (so to speak) *with* it – by its occurrence in S making (i.e., constituting) S's awareness of something (e.g., an external object). (Dretske 1997:6)

Higher-order theories describe state consciousness in the first way, as a mental state that is the intentional object of a higher-order thought or perception. Flat theories describe state consciousness in the second way, as a mental state that is an act by which the creature represents something in the world. Mental states are conscious when used to represent objects in a particular way, viz. consciously. Note that this description does not include what is particular about the way conscious states represent objects, as opposed to the way unconscious states represent objects. In a moment I will argue that this omission is a signif-

icant defect in Dretske's account of state consciousness. Nonetheless, Dretske's proposal that we think of conscious states in terms of relations between creature and world rather than relations between mental states is an important contribution.

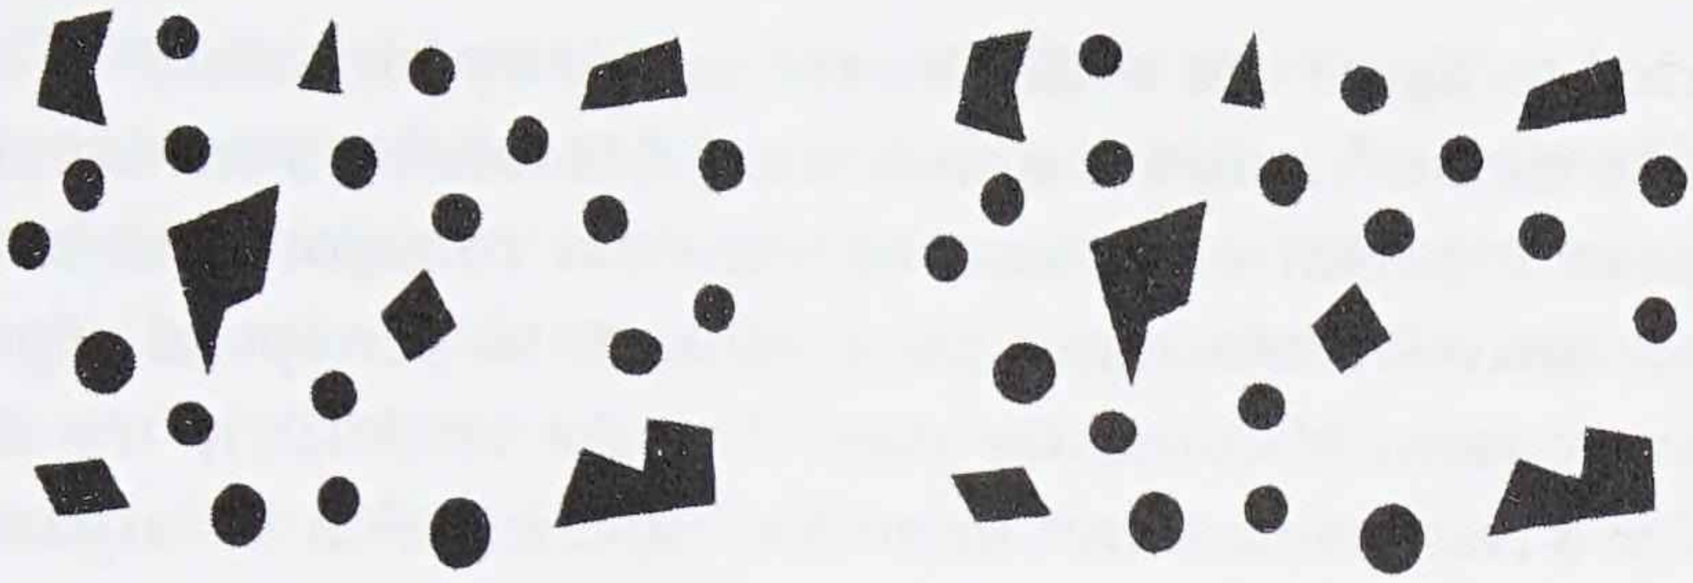
His commitment to a deep and pervasive externalism led Dretske this insight. In Dretske's view the content of any mental state is determined by creature-world relations, suitably described (Dretske 1995). The higher-order claim that state consciousness *implies* transitive consciousness of that mental state[72] assumes state consciousness is determined internal to the mind, an assumption directly counter to Dretske's externalism. Flat theories counter the higher-order assumption by arguing that conscious states are acts of the creature by which it represents something in the world. By an 'act of the creature' I mean that conscious states are a way the creature has of representing objects in the world. Seeing a tree (either consciously or unconsciously) is a way of representing the tree that differs from touching the tree. Similarly, consciously representing objects differs from unconsciously representing objects. They are different ways by which a creature represents things in the world. The difficult part in telling such a story, of course, will be in explaining the difference between conscious and unconscious ways of representing.

Unfortunately, Dretske does not clearly distinguish between conscious and unconscious states, let alone explain the difference. This ambiguity shows up in his use of the phrase 'conscious of'. In the following discussion I will begin by using the phrase without qualification, just as Dretske does. Then, once we have his argument in hand, we can see whether the argument is more forceful if Dretske means 'transitively conscious of' (or perhaps just 'represents') or 'state conscious of' (that is, 'has a conscious state that represents'). I will argue that the latter usage poses an important puzzle that higher-order theories of state consciousness have yet to solve.

2.1 Dretske's trouble spot

Briefly, Dretske's objection is that conscious states can differ, there can be a conscious difference between them, without consciousness *of* a difference. Therefore state consciousness should not be described as a state one is conscious of being in. To illustrate the point, Dretske instructs us to look at ('foveate' as he puts it) all parts of the following sets of objects, identified as Alpha and Beta (Figure 1).

What you probably did not notice is that Alpha contains a spot, called Spot (second from the right, lower row), that is not included in Beta. Because you



Alpha Beta

Figure 1. Dretske 1993b: 273

were conscious of Alpha and conscious of Beta and Alpha and Beta are different, you were conscious of different things. If you didn't happen to notice Spot, though, you were not conscious of seeing a difference (Dretske 1993b: 273). You had two different conscious experiences but were unaware of the difference.

Dretske explains how this case poses an objection to higher-order theories of consciousness:

> 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. E(Spot) [the experience of Spot] is conscious, and it constitutes a conscious difference between E(Alpha) and E(Beta) even though no one, including the person in whom it occurs, [is] conscious of it. (Dretske 1993b: 278)

The word 'experience' here is confusing because it is unclear how 'experience' relates to state consciousness. Dretske does not define 'experience', except to note that it is an example of a "concept-free mental state," as opposed to belief which is a "concept-charged state" (Dretske 1993b: 263). When perception is concept-free, one has a perceptual experience as opposed to a concept-charged perceptual belief. On Dretske's view, perceptual experiences entail consciousness of something in the world. Take the case of Spot. You look at Alpha and then at Beta. Spot is the difference but you don't notice Spot. Nonetheless, Dretske claims that your experience of Spot is conscious. Your experience of Spot is conscious presumably because it is a perceptual experience. You saw Spot, so according to Dretske, you were conscious of Spot. Dretske's first of four postulates on conscious experience states the entailment unequivocally:

(1) S sees (hears, etc.) x (or that P) $\Rightarrow$ S is conscious of x (that P)[73]

Now we need to figure out what Dretske means by 'conscious of' in his argument. If he means 'transitive consciousness', it follows from axiom (1) that any perceptual experience is a case of transitive consciousness of the thing experienced. Since Dretske limits his arguments to perceptual experience as a paradigmatic form of conscious state (Dretske 1995:103), 'transitive consciousness' is a reasonable substitution for 'experience' in his argument. And if this is what Dretske means by 'experience' we can see why he thinks this case disproves higher-order theories. If you see Spot, then you are transitively conscious of Spot (according to (1)). But you may not be transitively conscious of seeing Spot, so conscious states cannot be described as states you are transitively conscious of being in. You are not transitively conscious of being in the sensory state of seeing Spot, yet you are transitively conscious of Spot. So state consciousness must be something other than transitive consciousness of sensory (and other lower-order) states. Put this way, without the misleading word 'experience', it becomes clear how much of Dretske's argument rests on his axiom (1).

It also becomes clear how higher-order theorists can respond: Dretske is conflating transitive consciousness and state consciousness.[74] On a higher-order theory, transitive consciousness can occur in the absence of state consciousness – indeed, many unconscious intentional states are instances. By seeing Spot I may become transitively conscious of *Spot*, but only by having a higher-order state (a higher-order thought or perception) will I be transitively conscious of *seeing Spot*. Only by having a higher-order state will my sensory state of seeing Spot be a conscious state. When you look at Alpha and then at Beta, failing to notice Spot, your higher-order representations of the two sensory states are identical even though the sensory states themselves may differ. Whether your higher-order representation is either perception or thought, it fails to include a representation of seeing Spot. When Spot is pointed out, you become transitively conscious of seeing Spot by forming a higher-order representation of seeing Spot.

Given that so much of the argument against higher-order theory rests on (1), what has Dretske to say in defense of this claim? Surprisingly, nothing. Presented in the context of an argument for distinguishing between consciousness of things and consciousness of facts, Dretske states (1) without supporting justification. Of course perceptual states such as seeing, hearing and smelling certainly are types of conscious states, but they are also types of unconscious states, at least there seems good reason to believe perceptual states can be unconscious. Dretske's axiom (1) is as unfounded as the higher-order theorist

claim that state consciousness implies transitive consciousness of that state. Both claims are contentious and must be argued rather than assumed.

If Dretske is using 'conscious of' in the sense of 'transitive consciousness', then thoughts about a problem while doing some other task, repressed desires, and covert sensory processing of various sorts would all count as 'conscious.' Since Dretske often uses the term 'creature conscious' in describing what kind of consciousness he is trying to explain, a reader would be justified in concluding that 'conscious of' in (1) is meant this broadly. But if this is what Dretske means, then he has simply missed the point of higher-order explanations. Their target is an explanation of the difference between conscious states and unconscious states. By asserting that all perception is conscious, Dretske conflates the difference to be explained into a single category.

There is a second way to interpret (1) that might prove more effective. Dretske could say that sensation is always conscious because sensory terms – seeing, hearing, etc. – *mean* that the sensation is a conscious sensory state. On this reading, then, there could be no unconscious seeing because 'seeing' *means* 'conscious seeing.' Dretske's description of unconscious sensory states supports this interpretation of (1). For example, when discussing the function of state consciousness, Dretske asks, "Why aren't we all, in each sense modality, the equivalent of blindsighters who (it seems) get information about nearby objects (indicated by statistically significant performance toward those objects) needed to determine appropriate action without experiencing (seeing) them?" (Dretske 1995:119) Blindsighters acquire information about objects through their visual systems, but it seems Dretske would say this is not 'seeing.' Though the senses yield representational states on this interpretation of Dretske's account, they are not seeing, hearing, tasting, etc. when the process is unconscious.

On this second interpretation, Dretske means that when one sees (hears, etc.) x, then one is *state* conscious of x. Seeing something is only really seeing when one consciously sees, and conscious seeing implies being in a conscious state that represents the thing seen. This interpretation has the advantage of fitting neatly with a flat description of state consciousness as an act of the creature whereby the creature becomes state conscious of things in the world. A creature is only *state conscious of* something if it is in a conscious state that represents that thing. Dretske might intend his fourth axiom to make the same point (Dretske 1993b:270):

(4) S is conscious of x or that $P \Rightarrow$ S is in a conscious state of some sort.

Unfortunately, though, (4) yields the same two readings as (1). Either all intentional states are conscious states and Dretske simply fails to recognize the distinction between conscious states and unconscious states, or 'conscious of' must be read as 'state conscious of'. I believe the second reading fits more closely with the rest of Dretske's argument, although it has its own difficulties.

The problem with this use of 'seeing' as 'conscious seeing' (and 'hearing' as 'conscious hearing', etc.) is that it obscures the relations between unconscious and conscious states. The strong physical and functional similarities between unconscious and conscious sensory states suggest that both sorts of sensory states should be described as cases of 'seeing,' 'hearing,' etc. We can then clearly distinguish unconscious from conscious sensory states by appending the appropriate modifier 'unconscious seeing' or 'conscious seeing.' If we using 'seeing' exclusively for 'conscious seeing,' then how can we talk about 'unconscious seeing' without contradiction? To be sure, blindsight vision is odd, and we may be tempted to use scare quotes when we talk about what they can 'see.' But this temptation is grounded in the faulty view that all sensation must be conscious because it seems to us as if all our sensations are conscious. We now know that states of mind are not always what they seem. And the more we learn about neurological similarities between unconscious and conscious sensory processing, the harder it becomes to reserve sensory terms for conscious states only.[75]

Whether or not we call unconscious sensory states 'seeing' or 'hearing,' the advantage to the second interpretation is that it admits blindsight as a type of unconscious sensory state. Therefore, Spot might count against higher-order theorists after all. Dretske can argue that he is not conflating transitive consciousness and state consciousness, as the first reading suggests, because blindsight counts as a case of transitive consciousness without state consciousness. In blindsight, one is transitively conscious of something by virtue of an unconscious state. Blindsight patients have sensory states that represent objects but those sensory states are not conscious. Seeing Spot, however, is a case of state consciousness without transitive consciousness of those states. Though one is not transitively conscious of it, seeing Spot is a conscious state due to the following considerations. Alert people with fully functional visual systems look at Alpha and every part of Alpha. Spot is part of Alpha, so viewers look at – 'see' in the sense of 'consciously see' – Spot, even though they are not transitively conscious of seeing Spot. Therefore viewers have conscious states without being transitively conscious of those states. If this argument goes through, then on one reading Dretske's objection to higher-order theories holds. Seeing Spot is a conscious state even though no one is transitively conscious of seeing

Spot. So state consciousness cannot be explained as transitive consciousness of lower-order states.

2.2 Spot-sight and thimble-seeking

By contrasting the way we see Spot, call this Spot-sight, with the way blindsight patients 'see' stimuli in their blind field, Dretske's objection becomes much more persuasive. It is still not a knockout objection, however, because higher-order theorists can reply that Spot-sight is more like blindsight than conscious vision. The objection hinges on whether Spot-sight seems more like a conscious sensory state or an unconscious one. Thus, the case results in a stand off. In support of Dretske, we are not blindsight patients. We have fully functional visual systems which we use to look at Alpha, and the result of this task is a conscious state (or perhaps several). The question is whether the content of this conscious state includes Spot. If the conscious state is produced by looking at every part of Alpha as instructed, it seems reasonable that the conscious state includes a representation of every part of Alpha – in particular, it includes a representation of Spot.

In support of higher-order theory, on the other hand, seeing Spot is not like seeing Alpha. If asked, we would deny having seen Spot, and presumably we could not initiate any response to Spot in particular. Since it is hard to imagine any reason to respond specifically to Spot, consider the analogous situation presented in Daniel Dennett's game of Hide the Thimble. In the game, a thimble is placed in plain sight, such as on a bookshelf in a sea of knick-knacks. The object of the game is to find the thimble before the other players. When a player spies the thimble, she quietly sits down, trying not to give away any clues as to where she saw the thimble. The fun in the game comes when almost everyone is sitting down and one of the final players – 'Betsy' in Dennett's game – is looking right at the thimble but fails to notice it (Dennett 1991: 334).

As I read Dretske (the second interpretation), he would say that Betsy consciously sees every item on the bookshelf because her eyes are open and she is actively scanning every part of the bookshelf. We can even stipulate that she foveates the thimble, just to make the analogy with Spot-sight closer. Seeing the thimble is like seeing Spot – it is part of a conscious state even though Betsy is not state conscious of seeing the thimble in particular. A higher-order theorist would argue that the way Betsy sees the thimble is clearly different from the way she sees other items in the room, such as the bookshelf. If we asked, Betsy would no doubt report seeing the bookshelf. And her ability to respond to the bookshelf is evidenced by her behavior – scanning the shelves, not run-

ning into the corners. By contrast, Betsy would deny seeing the thimble and does not respond appropriately to the thimble. She does not sit down.

So seeing Spot or the thimble does not seem to be completely unconscious, like blindsight, nor completely conscious, like seeing Alpha or the bookshelf. Rosenthal deals with this troublesome middle-ground case in terms of general and specific representation. Before she noticed the thimble, Betsy had only a general representation of the bookshelf or the knick-knacks. When Betsy notices the thimble she acquires a more specific higher-order thought. She now represents the thimble *as* a thimble and responds to it as such (Rosenthal 1995:365). This distinction between general and specific forms of representation is promising, but within the context of higher-order thought theory, problems remain. It seems that Betsy shouldn't need to acquire a more specific thought if she is able to point to her sensations. Per hypothesis, Betsy senses the thimble, so she ought to be able to point to it, acquire the appropriate higher-order thought and thereby react accordingly. Looking at the situation another way, it seems Betsy already has specific thoughts about the thimble. Presumably she knows what a thimble is, at least sufficiently to identify one, so perhaps she needs to activate rather than acquire her thimble concept in this case. But if so, and her unconscious sensory representation is somehow able to activate Betsy's pre-existing thimble concept, it is unclear why it takes so long. Why isn't Betsy state conscious of the thimble at roughly the same time that she senses the thimble? It seems nothing is preventing her from wielding the more specific thought right away.

Finally, Betsy's ability to be state conscious of the thimble does not seem to be an ability to identify it *as* a thimble. It seems she could be state conscious of it even if she had no idea what thimbles are, as in the case of Spot. Dretske's test case assumes the reader has not noticed the difference between Alpha and Beta on first viewing and so is not state conscious of Spot in particular until it is later identified. But there may be very observant readers who do notice the difference upon first viewing and so are state conscious of Spot immediately. Must we require that these observers have a concept of Spot? Must they have a concept of Spot *as* Spot or *as* the difference between Alpha and Beta? If we do not set such requirements for sensation, why set them for state consciousness? Only if the best theory of state consciousness requires concepts should we accept this requirement for conscious states. Otherwise, as far as I can see, there is no justification for limiting the content of one's conscious states to one's concepts.

In this regard Dretske makes a useful distinction between thing-awareness and fact-awareness (Dretske 1993b:264–269). Awareness of things requires no

concepts of the thing perceived. Awareness of facts, on the other hand, requires the concepts necessary to express the fact of which one is aware. "If S is aware that *x* is *F*, then S has the concept *F* and uses (applies) it in his awareness of *x*" (Dretske 1993b: 265). In discussing the difference between thing-awareness and fact-awareness, Dretske describes studies in which monkeys are trained to respond to the larger of two rectangles. Eventually the monkeys learn to respond to the larger of any two differently-sized rectangles (A and B) presented. This achievement, according to Dretske, indicates the acquisition of the concept of the 'larger than' relation. He concludes, "we can say that they are aware of A, aware of B (thing-awareness), and aware that A is larger than B (fact-awareness)" (Dretske 1993b: 277). The critical point for Dretske (and for me) is that in learning situations like this, thing-awareness comes before fact-awareness. Thus, a person can be thing-aware of *x* without being fact-aware that *x* is *F*. Consider the monkeys again. Dretske asks:

> How shall we describe the monkeys' perceptual situation *before* they learned to abstract this relation? Did the rectangles *look* different to the monkeys?...We can imagine the difference in size to be as great as you please. They were not fact-aware of the difference, not aware that A is larger than B, to be sure. But that isn't the question. The question is: were they conscious of the condition of A and B that, so to speak, makes it true that A is larger than B?
>
> (Dretske 1993b: 277)

It seems obvious to me that they were. As Rosenthal himself admits, conscious sensory differentiation outruns conceptual resources (Rosenthal 1997: 742). So in Dretske's terms, Betsy can be thing-aware of the thimble without being fact-aware that it is a thimble, and one can be thing-aware of Spot without being fact-aware that it is Spot (or the difference between Alpha and Beta).[76]

But the distinction between thing-awareness and fact-awareness does not solve the problem of state consciousness, even though Dretske may think it does.[77] Narrowing now to sensory states, we need to understand the differences among the following states:

1. Betsy's sensory state when she is state conscious of the thimble *and* also notices and responds to it.
2. Betsy's sensory state when she is state conscious of the thimble but does not notice it.
3. The sensory state of the blindsight patient when he is not state conscious of the thimble but is able to respond to it (through forced-choice guesses).

I suggest that these differences are best understood in terms of *degrees* of sensory consciousness. The idea that there are degrees of sensory consciousness is widespread. We commonly speak of 'fading in and out' of conscious sensory states toward one end of the continuum and of 'heightened' conscious sensory states toward the other. These degrees of consciousness can be viewed as degrees of coordination, where a high degree of sensory consciousness involves well-coordinated representations and a low degree of sensory consciousness involves loosely coordinated representations.

Type 1 cases exhibit the highest degree of sensory consciousness – well-coordinated representations oriented toward a specific task. Here Betsy coordinates her visual representations of the thimble, distinguishing figure from ground, with her proprioceptive representations of body position to prepare to sit down. Betsy also needs to incorporate visual and auditory representations of her opponents' locations so as not to alert them to the location of the thimble. Type 2 cases exhibit a lower degree of sensory consciousness, involving a more loosely coordinated representation oriented toward general tasks. When Betsy begins her search for the thimble, her sensory representations are only loosely coordinated. The contents of her conscious sensory states are more random than in the Type 1 case: an auditory representation of a friend's giggle, a visual representation of the bookcase, an olfactory representation of cookies baking in the kitchen. Type 3 is a case of unconscious sensory representation; as such it is not coordinated into a conscious sensory representation at all. Some of Betsy's unconscious sensory states might include the kinesthetic representations of her legs as she walks or auditory representations of the grandfather clock ticking in the corner.

In the next chapter I will elaborate on these cases and propose a theory of sensory consciousness designed to account for the differences of degree among them.[78] In particular, considerably more needs to be said about the difference between Type 1 and Type 2 cases which are coordinated into conscious sensory representations, and Type 3 cases which are not so coordinated. Already my example of kinesthetic leg representations while walking shows that simple coordination will be insufficient, since walking seems a fairy clear instance of coordination that is nonetheless unconscious. As a hint, I will argue that the content of conscious sensory states has a temporal element that is lacking in unconscious states: conscious states represent the present moment. Before turning to this difficult point, I will conclude this chapter by noting why representationalists have difficulty accommodating degrees of consciousness.

The problem for a representationalist is that representations do not admit of degrees. Either something is represented or it is not. Despite this fact, there

are ways to accommodate degrees of consciousness without violating the Universality of Representation Thesis: the mind is constituted by representational properties and their functional relations.[79] Lycan, for example, argues persuasively in favor of degrees of state consciousness as a function of the number of inner monitors. In Lycan's view inner sensing is accomplished by 'ranks and ranks' of inner monitors dedicated to scanning different mental states. These mental states are conscious when an inner monitor scans them and produces a higher-order representation of them. More inner monitors mean more conscious states, hence a higher degree of state consciousness. Creatures with, say, just one monitor would have a very low degree of state consciousness, probably nothing comparable to our conscious states (Lycan 1996: 36–42).

But this account will not help us with Spot-sight and thimble-seeking. As we discovered earlier, higher-order monitors produce representations of sensory states, not things in the world. So Betsy's inner monitors scan her sensory states to provide information about how she is seeing. They do not scan the world for a thimble; that's the job of external senses. So higher-order perception theory cannot solve the problem of thimble-seeking and Spot-sight even though it provides an account of degrees of state consciousness. In the next chapter we will see if a second sense theory can do better.

3. Conclusion

Higher-order theories have proven problematic on several grounds. In order to have higher-order thoughts about novel sensory states, some form of demonstrative reference is required. But in the absence of a sensory mechanism it is unclear how such a demonstrative reference could be secured. Higher-order perception theory can accommodate novel conscious sensory states but seems to offer an explanation of introspection rather than state consciousness. These problems are highlighted in the case of Spot-sight where we seem to be state conscious of Spot though we are not state conscious of Spot in particular. To deal with Spot-sight we need a theory that explains how the conscious sensory state that represents Spot as an undistinguished part of Alpha differs from the conscious sensory state that represents Spot in particular. The elaboration of such an explanation is the task of the next chapter.

Chapter 3

Solving the problem of Spot-sight

At the end of the last chapter we were faced with a puzzle we now need to resolve. In thinking about Dretske's case of Spot-sight and Dennett's game of thimble-seeking, we found three basic types of sensory states that a theory of sensory consciousness needs to explain.[80]

1. being state conscious of x[81] and noticing/responding to x
2. being state conscious of x without noticing x
3. not being state conscious of x but nonetheless able to respond to x

Higher-order theories have trouble distinguishing between Type 1 and Type 2 cases. According to higher-order theory, a state is conscious when it is represented by a higher-order state, either a thought or sensation. But Type 2 cases are examples of conscious sensory states without higher-order representation. We are state conscious of Spot when we consciously look at all parts of Alpha, but we do not notice Spot, we do not have a higher-order representation of seeing Spot. In assiduously searching the room for the thimble, Betsy consciously sees the thimble on the bookcase, but she does not notice the thimble, she does not have a higher-order representation of seeing the thimble. On a higher-order theory, one cannot have conscious states unless one has higher-order representations; one cannot be state conscious of anything unless one represents seeing (or otherwise representing) it. Therefore, higher-order theory must deny that Type 2 cases are cases of conscious sensory states.

On the other hand, Dretske has trouble distinguishing between Type 2 and Type 3 cases. On the interpretation of Dretske that poses Type 2 cases as a problem for higher-order theory – all seeing is conscious seeing – he is left without a clear distinction between consciously seeing x and unconsciously seeing x.[82] He therefore cannot explain how blindsight patients are able to respond to visual stimuli, since they must represent these stimuli in some way in order to respond to them. Nor can he explain why Betsy cannot respond to the thimble, since she consciously sees it in Dretske's view.

In both theories, Type 2 cases cause a problem: they are neither full-fledged conscious sensory states nor are they unconscious. Type 2 cases somehow fall

between paradigmatically unconscious and paradigmatically conscious sensory states. So how should we characterize the 'between-ness' of Type 2? I suggest it is one of degree: Type 2 cases are conscious sensory states, but they exhibit a lower degree of sensory consciousness than Type 3 cases. The first task of this chapter is to elucidate this notion of degree and to account for it within a theory of sensory consciousness. Part 1 proposes that degrees of conscious sensory states are best understood as degrees of coordination necessary for accomplishing the kind of task at hand. Briefly, Type 1 cases involve highly coordinated sensory representations necessary for specific tasks, and Type 2 cases involve loosely coordinated sensory representations necessary for more general tasks.

The coordination hypothesis is insufficient to explain the difference between Type 1 and Type 3 states, however, as some tasks involve highly coordinated representations that are nonetheless unconscious. Think, for example, of typing or reading. While conscious states are certainly involved in these activities, the detailed sensory-motor representations involved in these tasks is arguably unconscious.[83] In order to account for the difference between conscious and unconscious coordination, Part 2 argues that conscious states are coordinated sensory representations endowed with an additional representational content: conscious states represent the present moment. The mechanism that endows this peculiar content is the second sense. As Part 3 explains, the job of the second sense is to select sensory representations and coordinate them into a representation of the world 'now'. In conclusion, Part 4 reviews the case of Spot-sight in light of the second sense theory to see if we have found an adequate account of all three types of state to be explained.

3.1 Coordinating sensory consciousness

The first step toward this goal is to clarify the suggestion that Type 2 cases exhibit a lower degree of sensory consciousness than Type 1 cases. The degree of coordination among sensory representations is determined by the sort of task undertaken. When Betsy is just beginning her search, her sensory representations are loosely coordinated. When Betsy narrows her search to focus specifically on thimble-sized objects, her sensory representations become more well-coordinated. Before explaining this distinction between 'loosely coordinated' and 'well-coordinated', let me underscore the point that this continuum is only the beginning of an account of sensory consciousness. For the moment we are looking at a possible explanation of the distinction between Type 1 and Type 2

cases. Once we establish that degree of coordination is successful in explaining this distinction, we will still need to explain how Type 1 and Type 2 cases differ from Type 3 cases.

In Type 1 cases one is state conscious of and also notices and responds appropriately to the stimulus. In such cases, sensory representations are coordinated to accomplish a specific task. Noticing the thimble, for example, requires that Betsy form a representation that coordinates shape, size and color representations about a specific target object, the thimble, and excludes distractor representations or surrounding objects and features. Type 1 cases are the sort that convince us sensory consciousness must have an important function because it seems that if we didn't have a conscious sensory state about the object right *there* in front of us we would have no hope of effectively responding to it. Only after multiple psychological studies have shown how much mental work is done by unconscious processes does it seem possible that quite detailed tasks such as typing might be accomplished without equally detailed Type 1 sensory consciousness.

Type 1 cases also seem to support the cognitivist idea that sensory consciousness requires some sort of conceptual grasp of the objects represented by conscious sensory states. As Dennett puts it, some kind of 'microcognition' is required for sensation to be more than mere visual machinery on the order of a camcorder (Dennett 1994:512f). Since sensory consciousness is a form of sensory state, it too ought to require some kind of cognitive uptake to be more than mere machinery. I am willing to grant that in Type 1 cases, concepts are likely to be involved. Our ability to form and wield concepts is too useful not to be exercised in successful response to specific tasks. But I disagree with Dennett's claim that only cognitive uptake will allow us to distinguish true vision from a camcorder recording.[84] One reason to oppose Dennett's cognitivism is that it cannot account for Type 2 cases. In Type 2 cases one is state conscious of but not able to respond to the stimulus. On the cognitivist view, if one is state conscious of *x* then one has a concept of *x*, so why no response? One answer is to deny that one is state conscious of *x*. Rosenthal, for example, explains Type 2 cases in terms of insufficiently specific concepts. While Betsy is state conscious of the bookshelf, of which the thimble is a part, she is not state conscious of the thimble in particular because she has not yet acquired the more specific higher-order thought about the thimble. When she acquires the appropriate higher-order thought, then she becomes state conscious of the thimble.[85]

I think Rosenthal is on the right track here in distinguishing Type 1 and Type 2 cases as specific and general, respectively. But as I noted in Chapter 2, his description of the distinction in terms of concepts is problematic. Rather

than draw the distinction in terms of the specificity of concepts, I suggest we draw the distinction in terms of the specificity of task: highly coordinated representations are necessary for accomplishing a specific task, whereas loosely coordinated representations are necessary for performing general tasks. At the start of the game Betsy scans the room in hopes of spying the thimble. But the room is too full of stimuli, so her scanning efforts overload her with a fairly random collection of sensory information. Her task is general, characterized by broad directives such as 'get spatially oriented' or 'figure out what to do next.' General tasks can also be described, logically, as the opposite of specific ones: task that are not specific or are prerequisite to engaging in a specific task. In Betsy's case, she must become familiar with the spatial layout of the room – the walls, the location and shape of the furniture, the position of onlookers and competitors – before she can dispense with the task of learning the overall room features and focus on specific thimble candidates.

The distinction here is not simply one of size, although size is one factor involved. Large objects, such as a bookshelf, include smaller objects and so a general categorization of room features is likely to include large objects. Equally important are visual categories. A pitcher set among a collection of pitchers is less easily discriminated than the same pitcher standing alone on a buffet. The solitary pitcher may be among the objects Betsy discriminates on her first sweep of the room, while discriminating a single pitcher amidst a collection would require more careful attention. Task specificity is a function of factors such as the size, uniqueness and mereological relations of the target. Note that effort is not one of the factors that determines task specificity. Betsy is working just as hard when she is scanning the room as when she is focussing on particular objects. Degree of effort is, I would imagine, highly correlated with degree of coordination. The former does not determine the latter, however. A general overview of the room will reveal any very large, unusual object that is not easily perceived as a part of another object. A much more specific search will be required to identify a small object, one that is similar to its neighbors or that can be seen as part of a larger object. For this sort of task, Betsy must be able to select and coordinate just those sensory representations that meet very specific criteria.

It may seem that highly coordinated representations are also needed in the general task, since Betsy must be doing some binding operations to represent the shape of the room and the position of bookcase, chairs, and so on. This is a good point to remember that we are considering two forms of *conscious* sensory states. The claim is that Type 1 and Type 2 states should be distinguished in terms of degree of representational coordination. Type 1 states involve highly

coordinated representations in order to perform specific tasks, whereas Type 2 states involve loosely coordinated states in order to perform general tasks. So, while Betsy's representations of various room features must be coordinated to give Betsy the topographical framework with which to narrow her search pattern, this coordinated representation need not be *conscious*. There need be no conscious state that includes a representation of the relevant features of the room at once. This distinction will become clearer after I have said which sorts of representations are included in conscious sensory representations and which are not. For the moment, imagine Betsy as she moves around the room, consciously representing a schematic tableaux of edges and shapes here, another over there, and so on, in successive moments. These loose, impressionistic schemas do need to be integrated into a representation of the layout of the room, but this integration need not be conscious.

So how is this an improvement on the idea that concepts are necessary for conscious sensory states? Isn't the 'coordination of sensory representations' simply a matter of forming conceptual representations of objects? Not necessarily. As I see it, coordination of sensory representations is a precursor to conceptual representation. By 'concept' here I mean simply the ability to reidentify an object, to recognize the same thimble again.[86] Coordinating disparate representations about an object facilitates reidentification, but coordination and reidentification are separate abilities. A person may have a well-coordinated sensory representation of a thing and yet have no concept of what it is. Consider a desert nomad wandering into a deep forest after a lifetime spent surrounded by sand and stone. He would surely be able to *see* individual trees as he wandered about, but he would not see *that* they are trees, he would not know they are trees, at least not purely by acquaintance. The coordination of representations certainly feeds the wheels of concept production, but they are not identical. Coordination is a fleeting thing – what is well-coordinated one moment may be only loosely coordinated (or completely disparate) another. Houses of cards and picture puzzles attest to the temporary nature of coordination. Conceptual abilities, on the other hand, are more lasting if not indelible. If I have a concept today, I will have it tomorrow so long as I continue to make use of it.[87] As we will see, however, coordination of sensory representation is a present-tense operation.

Because coordination of sensory representation is necessary for accomplishing a given sort of task, it is reasonable to assume that concepts will be pressed into service wherever useful. I am in no way suggesting that concepts never affect the content of conscious sensory states. My point is that concepts are not necessary for coordination of sensory representation, and so on my ac-

count, are not necessary for sensory consciousness. In other words, a creature without any concepts whatsoever could have conscious sensory states. Sensory consciousness is a matter of the coordination of sensory representations, either to a high degree as when a specific task is at hand or a low degree as when a more general task is underway. Coordination of sensory representations into conscious sensory states facilitates concept acquisition and often makes use of existing concepts, yet there is no conceptual requirement for the production of sensory consciousness.[88]

Two questions about degree of coordination need to be addressed before going on to look at the distinction between conscious and unconscious sensory states. The first question deals with the number of sensory states necessary for coordination, and the second concerns the degree of coordination necessary for sensory consciousness. First, we can ask how many sensory states must be coordinated to constitute sensory consciousness. As the term is applied to us, 'sensory consciousness' involves multitudes of features represented. Normally our conscious sensory states are cluttered with many varieties of sensory representations, the smell and taste of fresh coffee, the feel of the smooth weight of the cup, the sound of a violin concerto. So, on a first pass, it seems a sensory state that contains just one feature would not be at all like our sensory conscious states.

On closer investigation, we can make the point even stronger. I find it impossible to imagine a conscious sensory state that represents a single feature. If that feature were the color white, then the display would have to be absolutely homogeneous, with no wrinkles or shadings. For if these occurred, then there would be two features, two shades of white, represented rather than one. Furthermore, all of a person's proprioceptive receptors would have to be stilled, since in the absence of other forms of stimulation, these normally unconscious sensory states would fill the vacuum of conscious content. Apparently sensory deprivation tanks result in just this effect. In the absence of normal sensory input, each subtle breath or movement is amplified far beyond its usual range.[89] So, to represent a single feature would require eliminating any cues that could distinguish subject from object; no feature externally or internally could serve to make this distinction. All that would exist would be a conscious representation of white, with no subject. Such a notion makes no sense. At least as many sensory states must be coordinated to provide a distinction between subject and object.

Less hypothetically, it may seem that mystics can achieve states of pure consciousness ruled out by the above considerations. The practice of deep meditation is supposed to result in just the dissolution of subject and object dismissed

as nonsensical. In response, let me clarify my point a bit. As I said, it is nonsense to claim that a conscious representation of white could exist without a subject. For there to be a representation of something else, there must in some sense be two things: the representation and the represented. The represented could fail to exist; it could be merely intentional. But there must be some object or feature the representation purports to represent. This sort of directedness is an essential relational feature of representations. As far as I understand mystical states, the goal is to eliminate such relational aspects of existence so as to reveal the true nature of all as one. Thus, mystical states cannot be representational states, so they cannot be conscious representational states. Mystical states are reasonably considered to be some form of consciousness, but they are not conscious sensory states in the sense under investigation.

As for the second dimension of degree, I have said that conscious sensory states are *coordinated* sensory representations of the present moment. Coordination is a matter of degree, so how much coordination is required? How well integrated must the sensory states be in order to count as conscious? When well-coordinated, conscious sensory states involve specific tasks, such as when Betsy is seeking the thimble. When more loosely coordinated, conscious sensory states involve more general tasks, such as avoiding obstructions. So, how loose can conscious sensory states get? What is the minimum amount of coordination necessary to constitute sensory consciousness? It seems to me that degree of coordination parallels degree of sensory consciousness, such that a little coordination would constitute a little sensory consciousness. Those in-between states as we drift off to sleep, characterized by a confused, random mix of sensations, are conscious sensory states on the edge of unconsciousness. Such loosely coordinated states can barely be considered task-oriented at all, involving only the most general action like changing position or pulling a blanket around your head. On the other end of the continuum, the most well-coordinated states involve highly focussed attention on a very specific task, such as the jeweler requires when repairing the intricate inner workings of an antique watch. Most of our conscious sensory states fall between these extremes, shifting from more to less coordinated and back again as we take up one task and leave another.

Thus we can distinguish Type 1 and Type 2 cases of sensory consciousness by postulating a continuum of coordination from highly coordinated sensory representations necessary for a specific task to loosely coordinated sensory representations required for more general tasks. Now what about Type 3 cases? Blindsight presents one dramatic example of the way an unconscious stimulus can facilitate a fairly specific task. But we don't need to consult neuroscience

literature for examples of specific tasks that are accomplished by highly coordinated, yet unconscious, representations. While sitting here typing, scores of minute tactile and motor representations of finger positions and movement sequences are needed to produce the stream of words on the keyboard, although none of these representations reach sensory consciousness. Why are all these processes unconscious? What is the critical distinction between unconscious and conscious sensory states? It can't simply be coordination, because some activities such as typing require coordinated sensory representations which are clearly unconscious.

Let's go back to the second of my examples of the difference between unconscious and conscious sensory states:

> For the last 20 minutes I have been shifting around in my chair, crossing one leg and then the other, sitting forward, then back. Only now when I have turned my attention away from my task do I become aware of the bodily condition that has been causing my movement. Now I notice the ache in my shoulders and the cramped feeling in my legs.

One minute my leg and shoulder cramps instigate their effects without benefit of sensory consciousness, and the next minute they are conscious sensory states. What is the difference? When unconsciously representing the tension and inflammation in my legs and shoulders, I respond to the pain signals without incorporating the information about my bodily condition into my overall representation of what is happening at the moment. Though my sensory representations of bodily damage are sufficiently coordinated with motor systems to produce my periodic fidgeting in the chair, they are not coordinated into this broader representation of what is happening now. My writing task encompasses my world at that moment, and anything else is represented as a peripheral element in that world or not represented as part of that world at all. When the task is finished and my attention is free to turn to other concerns, my representation of leg and shoulder cramps is then incorporated into my coordinated sensory representation of what is happening at the moment.

I suggest that it is this overall representation of what is happening at the moment that constitutes the content of conscious sensory states. Sensory representations coordinated into conscious states are the best approximation of what the world is like at the present moment. In the rest of this part I will unpack the two central elements in this formulation: (1) that they are the *best approximation* of the world, and (2) that they are representations of the world *at the present moment*. The final section then considers the role of a representation of the world at the present moment in *decision and action*.

3.1.A The best approximation of the world

First, conscious sensory states are the *best approximation* of what the world is like at the present moment. The coordination processes involved in sensory consciousness are, on my view, aimed toward tasks. The sensory representations coordinated into conscious states are selected on the basis of their usefulness in accomplishing the task at hand. When only general tasks are undertaken, with no particular end in view, it may be best to incorporate whatever sensory representation happens to be floating about. Often when we are in between specific tasks our conscious sensory states are filled with such random representations. Body pains, a glimpse of the mail carrier on the porch, and the taste from a swig of hot coffee are among the sensory representations that flood the mind while the decision about the next action is deliberated. Or, as in the case of Betsy, a certain amount of general sensory information may need to be acquired before more specific focus is possible. Betsy must learn the lay of the land before she can turn to more specific tasks. At that point, sensory representations are more carefully selected, combined or even confabulated in order to form a coherent representation of the thimble.

The so-called 'cutaneous rabbit' experiments described by Daniel Dennett and Marcel Kinsbourne show how subjects often misrepresent stimuli in order to form coherent representations. Subjects in these studies receive three taps on the wrist then three taps on the elbow and finally three taps on the upper arm. Strangely, subjects do not report feeling these three sets of three taps but instead report feeling a series of taps evenly spaced along the arm – as if a rabbit is hopping from wrist to shoulder. However, if subjects receive only the first three wrist taps without the subsequent elbow and upper arm taps, they then accurately report feeling all the taps at the wrist – the 'rabbit' hops in place (Dennett & Kinsbourne 1992:186). One way to account for this phenomenon is to postulate a time delay required for the wrist-taps to be represented in a conscious sensory state, allowing later taps to influence the representation of earlier taps. The influence of later taps thereby causes a revision in the conscious representation of earlier taps. Dennett and Kinsbourne call this Stalinesque[90] editing because it involves manufacturing a representation of evenly spaced taps to appear in conscious sensory states.

But another possible explanation of the cutaneous rabbit reports is that the editing occurred *after* the taps were incorporated into sensory consciousness but before the final report (Dennett & Kinsbourne 1992 190f). Dennett and Kinsbourne call this second alternative Orwellian[91] editing because the subject's memory of a veridical conscious representation is changed. The subject

has fleeting conscious sensory states accurately representing three taps at the wrist, three at elbow and three at upper arm, but those conscious representations are subsequently wiped out and replaced by a report of evenly spaced taps along the arm. According to Dennett and Kinsbourne, there is no functional difference between Stalinesque and Orwellian forms of editing because there is no way to determine such a thing as a "charmed circle of consciousness" where all and only conscious sensory representations are located (Dennett & Kinsbourne 1992: 193).

No doubt significant methodological problems stand in the way of arbitrating between Stalinesque and Orwellian forms of explanation. Sometimes it seems as if the apparently hopeless task of verifying the physical substrate of conscious sensory states is the reason Dennett and Kinsbourne reject the difference between these two interpretations. The difference, if one can be determined, is a matter of exactly which representations are part of a conscious sensory state, whether the editing from three 3-tap intervals to evenly spaced taps occurs before or after the taps are consciously represented. Subjective reports cannot arbitrate between Stalinesque and Orwellian explanations because both explanations predict the same sorts of report. Therefore the only way to determine a difference here is to determine a *physical* difference between conscious sensory representations and unconscious ones. At this point we are still struggling to come up with a good operational definition of sensory consciousness and a long, long way from the sort of neurological identification of conscious sensory states that might provide independent verification of their contents. Still, there seems to be no principled reason why future research could not make some such identification and thereby distinguish Stalinesque from Orwellian forms of editing.[92]

Nonetheless Dennett and Kinsbourne strongly maintain that Stalinesque and Orwellian editing are *metaphysically* indistinguishable.

> Both the Orwellian and the Stalinesque version … can deftly account for *all* the data – not just the data we already have, but the data we can imagine getting in the future. (Dennett & Kinsbourne 1992: 193)

> … if one wants to settle on some moment of processing in the brain as the moment of consciousness, this has to be arbitrary. One can always "draw a line" in the stream of processing in the brain, but there are no functional differences that could motivate declaring all prior stages and revisions unconscious or preconscious adjustments. (Dennett & Kinsbourne 1992: 194)

There can be no functional property that is necessary and sufficient for a sensory state to be conscious, Dennett and Kinsbourne claim, because "there is no

further functional or neurophysiological property *K* over and above the properties that account for the various 'bindings' and effects on memory, speech, and other behavior" (Dennett & Kinsbourne 1992: 236). I agree with this statement, but do not see how it implies that there is no property *K* which could distinguish Stalinesque and Orwellian forms of editing. This follows only if one holds, as Dennett and Kinsbourne do, that effects on memory, speech and behavior *constitute* consciousness. Otherwise, one could envision any number of forms of binding that might serve as property *K*. In particular, I propose a form of binding that is necessary and sufficient for a sensory state to be conscious. A sensory state is conscious, on my account, when it is bound into a coordinated representation of the world at the present moment. There is a lot of work to be done in order to support this account and even more to determine a physical substrate for such coordinated representations. The question for Dennett and Kinsbourne, however, is why this sort of binding could not in principle account for sensory consciousness.

3.1.B Representing 'now'

What moment is represented by conscious sensory states as the present moment? As the cutaneous rabbit experiments show, the time represented is not just the time at which the sensory representations occur. Otherwise the wrist taps would be represented as grouped into sets, since the sensory representations occur in sets. This point is critical to understanding the distinction between *now representing* and *representing now*. On my view, all sensations are cases of *now representing*. This is just to say that all sensory representations are occurrent. There may be mental states such as beliefs that are non-occurrent representations; perhaps they are dispositional or have some other sort of structural, but inactive, ontological status. Sensory representations are not like this, however. Being active is an essential feature of sensory representation, and this feature accounts for the fact that sensation but not cognition requires the presence of the stimulus. So all sensations now represent some feature of the world. But on the proposed account only conscious sensations represent features *as* now. A representation of the present moment constitutes a different form of temporal coordination than one generated simply by stimulus onset. As argued above, sensory representations are coordinated so as to form the best approximation of the world at the present moment. In some cases, as with the cutaneous rabbit, this coordinated representation will involve confabulation. In all cases, this coordinated representation will involve the se-

lection of sensory representations according to some kind of temporal ordering principle.

To get an idea of how a representation of the present moment can differ from the time represented, consider the analogy of the stockmarket trade board. The board is continuously updated as stockbrokers watch various holdings jump up or down with market fluctuations. By watching the board, brokers can make quick comparisons among stocks and shout 'buy' or 'sell' as appropriate. The trade board is a relatively simple coordinating device, collecting together information about stock sales activity into a current market price. No choices need to be made about what sorts of information will appear on the board or when it will appear since all such mechanisms are determined in advance by the program designers. Yet the trade board illustrates how the time of representation differs from the time represented. At any moment the trade board represents the current price of stocks. If you log a sale while that price is on the board, that is the price you get. But other people are simultaneously buying and selling around the world which means by the time the sale is logged, another price is the *true* current price at that moment. The true current price is the actual value of the stock at that moment, which includes all the selling and buying as it occurs. But the true current price is impossible to represent because it takes time to collect all the necessary sales information. The price on the board represents the coordination of all the available information about stock activity when the price was listed on the board. This compilation is represented as the current market price. Since time marches on, however, the true current price has changed by the time the price represented as current is listed. Given the time required to produce a price for the trade board, the representation of stock prices that appears on the board *now* (and so by stipulation of trading practice represents them *as* current) actually refers to the price of stocks from a moment ago. The time of stock price representing (now) differs from the time represented by the stock price (a moment ago).

One might object that the rules of indexical reference stipulate against the above analysis. If 'now' must refer to the moment of utterance, it can only be used to refer to the present time. Dennett and Kinsbourne make this point in response to the suggestion of a "pink-elephant-now" thought as a test case for measuring conscious time. Because 'now' is indexical, "*It refers to its own vehicle of representation*, and hence bridges the gap between temporal properties represented and temporal properties of representings" (Dennett & Kinsbourne 1992:237). But common usage of 'now' is not so restrictive. Answering machines commonly state that the speaker is "not here right now," referring to the future time of the call rather than the present time of the utterance. The nar-

rator of a documentary may refer to the past events depicted in a segment of footage saying, "Now comes the decisive moment where the treaty is signed." In this case 'now' refers neither to the moment of utterance, nor to the moment of hearing the utterance; it refers to the events represented by the footage. The temporal elasticity of 'now' is most clear in the usage of a parent, where 'now' can refer to a range of time frames from 'this very minute' to 'sometime today' as individuated by context and tone of voice. Certainly some aspect of the present tense is integral to these various ways of using 'now', but reference to the moment of utterance is not a necessary element. In the case of the trade board, there is a representation of the present moment that refers to events in the past; the trade board represents the price of stocks as current, and refers to the price of stocks a moment ago. Dennett and Kinsbourne are correct to note that the "evidence concerns only the temporal properties *represented* in the thought . . . not the temporal properties of the representing" (Dennett & Kinsbourne 1992:237). My claim is that 'now' is part of the content, the representational glue that binds disparate bits of the world into a unified sense of how things are at a moment. But the moment represented need not be the moment of representing. On the contrary, I believe content systematically lags behind vehicle due to the time required for signal relay and coordination processes.

The lag between the time represented and the time of representation puts the stock broker in a tricky position. Because the stock price represented by the board as current is actually the price of a moment ago, the shrewd broker must anticipate the market trends in order to take advantage of prices listed as current. The experience of the broker allows her to predict whether a price will keep going up or whether the stock has reached a peak, and this experience is why brokers get paid for their services. Likewise, action on the part of conscious creatures involves predicting how the world will be at the next second, as well as representing the world at the present moment. As with the trade board, the moment represented as present by a conscious sensory state has already passed. It takes time for stimuli from whatever object or event is represented to reach the sensory organs and more time for sensory processing prior to the coordination of representations into a conscious sensory state. Therefore, effective action requires predicting how the world will be at the next moment in addition to representing some moment as present. This pull toward the future may be one of the reasons it is so difficult to determine the difference between Stalinesque and Orwellian forms of representational editing. In the cutaneous rabbit experiment, are the tap-representations edited *before* conscious sensory states are produced in order to provide the best approximation of what the world is like at the moment the taps occurred (Stalinesque), or are they edited

after conscious sensory states are produced in order to better anticipate the bunny's next move (Orwellian)? At this point in our understanding of sensory consciousness, the answer could go either way. *Pace* Dennett and Kinsbourne, however, there is a principled difference between the two answers. If the theory I propose is correct, there is a fact of the matter about which of these explanations is true. Either evenly spaced taps are incorporated into a coordinated sensory representation of the present moment or they are not.

It is important to note that neither Stalinesque nor Orwellian forms of editing require that the representation of time involves timed representations. Dennett and Kinsbourne are absolutely right to point out the problems with this sort of vehicle/content confusion (Dennett & Kinsbourne 1992: 188f). For example in the cutaneous rabbit experiments, it is not necessary that subjects produce an evenly spaced sequence of tap representations in order to have a representation of evenly spaced taps. This point is fairly easily accommodated by the Orwellian form of editing, since *per hypothesis* the revision occurs after the moment of consciousness. Thus there is no impulse in this case to 'fill in' extra tap representations to display in the theater of the mind. But even if the Stalinesque form of representational editing turns out to be true, there are at least two ways such a representation might be produced, both compatible with my proposal. One possibility is that very short intervals of time are compressed into a single conscious representation of 'now'. In this case, the evenly spaced taps would be represented as a unit of movement occurring at that moment. It makes good sense that the brain would interpret common patterns of movement holistically rather than represent each microsecond separately. Temporal clumping, like spatial clumping, would be an efficient short-cut in figuring out what sort of event/object is out there as well as in anticipating what sorts of changes are likely to follow. According to Richard Warren in his commentary on the Dennett and Kinsbourne article, the most evolutionarily plausible hypothesis of the cutaneous rabbit stimuli is that a single agent produced all the wrist taps, and this hypothesis would result in the representation of an evenly spaced sequence of taps. "This agent cannot readily jump abruptly from one of the three stimulated positions to the next within a single inter-tap interval (which ranged from 50 to 200 ms)" (Warren 1992: 231). Warren argues that a series of brief events such as the wrist tap series is likely represented as a 'temporal compound.'

A second Stalinesque possibility is that an evenly spaced sequence of conscious tap representations is indeed produced. Though a sequence is not *necessarily* represented by a sequence, sometimes it is best represented as such. Dennett and Kinsbourne acknowledge this point when they say: "If someone

thinks the thought, 'One, two, three, four, five,' his thinking 'one' occurs before his thinking 'two' and so forth. The example does illustrate a thesis that is true in general and does indeed seem unexceptioned, so long as we restrict our attention to psychological phenomena of 'ordinary,' macroscopic duration" (Dennett & Kinsbourne 1992:200). Because of the mistakes we make in representing temporal sequences of very short durations it is unclear exactly how those sequences are represented. If the Stalinesque interpretation of the cutaneous rabbit experiment is true, then a conscious representation of evenly spaced taps could be represented by either a single representation of sequence or a sequence of representations. The point is that neither Stalinesque nor Orwellian interpretations of the cutanteous rabbit experiment are committed to any form of vehicle/content confusion. We do not yet know how the rabbit stimuli are consciously represented, but we should not assume that the absence of knowledge entails the absence of a matter of fact.

On my view, whether the wrist taps are consciously represented as evenly spaced or not depends on whether they are appropriately coordinated into a representation of the present moment. Though I have some ideas about how a brain might accomplish the necessary coordination (see the Appendix), I have no idea how we might determine the content of any specific conscious representation simply by looking at the brain. Nonetheless, it is a matter of fact whether Stalinesque or Orwellian editing occurs, even if we can not determine which is the case due to methodological constraints.

What may be unsettling about this claim is the prospect that we could be mistaken about our own conscious sensory states. If, for example, the Orwellian description of the cutaneous rabbit phenomenon is true, then we have a veridical conscious representation of separate sets of taps at wrist, elbow and shoulder. Then, milliseconds later we report our conscious representation as having been a sequence of evenly spaced taps. How could we possibly be so wrong about our own conscious sensory states so soon after they have occurred? For Dennett, admitting no facts about conscious representation outside of subjective reports is better than admitting the possibility of subjective error.[93] Unsettling as this prospect may be, we should not be surprised that even this aspect of the much-maligned notion of introspective infallibility should fail. Because there is no necessary connection between the content of conscious sensory states and reports about that content, reports about one's own conscious sensory states, like any other reports, are fallible. Therefore, we cannot rule out the possibility that the Orwellian interpretation of our conscious sensory representations might be true, as disturbing as this possibility might seem.

In happier news, the proposal that conscious sensory states are coordinated representations of the present moment provides support for one of our other long-held ideas about sensory consciousness: the notion of a stream of consciousness. From the subjective point of view there seems to be a stream of conscious sensory states flowing from one moment to the next. This stream may be relatively coherent, as when focusing on a task, or it may contain a random collection of junk, as when one has no aim in particular. Whatever the content, each moment of sensory consciousness seems to form a unified collection, conjoined with moments preceding and following. The description of conscious sensory states as coordinated representations of the present moment dissolves the apparent tension between distributed, parallel brain processing and the unified representation of the world by conscious sensory states. Conscious sensory states represent the world *as* unified, even though many brain processes, perhaps even including the processes that compose conscious sensory states, are distributed and parallel.[94]

Let me expand a bit here, as the proposal of unified sensory representations may lead one to think that these representations must come together at some single place in the brain. The coordination of representations into conscious sensory states is determined by what things are represented as occurring now, and there is no reason to require this coordination occurs at a single neurological locus. Quite the contrary, in fact. On the present hypothesis, conscious sensory states are composed of unconscious sensory representations that are coordinated by some means so as to represent the present moment. If such a coordination were a matter of physically reproducing each sensory representation into one large, master conscious representation at a single spot in the brain, then one would expect there to be a single spot that housed all and only conscious sensory representations. No such spot exists. Various sorts of brain damage result in various sorts of deficits in conscious representation, but there is no form of brain damage that wipes out all conscious sensory representation while preserving other mental functions. So perhaps there is one master representation for each conscious sensory state, but the master representations occur at different points in the brain. One conscious state might be located in the anterior cingulate gyrus, for example, and another in the superior temporal sulcus. Such a proposal seems biologically arbitrary, however, with conscious sensory states popping up all over the brain like popcorn.

Pragmatic reasons, moreover, militate against any form of the single spot hypothesis. Reproducing sensory representations into one, master conscious sensory representation is redundant. If there are already representations lying around in the brain, why not just use them? In addition to the waste of redun-

dancy, the extra step of reproduction takes time. Simply using existing representations without reproducing them into a master representation is already a time-consuming operation. To go to the additional trouble of reproducing representations would be startlingly inefficient in the split-second world of conscious sensory representation. It would be more effective for the brain to simply make use of existing representations wherever they are located, which by current accounts, seems to be widely distributed across the cortex.[95] So we can set aside another of Dennett's worries – that coordination of representation will necessitate a single place where the representations 'all come together' (Dennett 1991:107). Representations do come together in the sense that their content is coordinated. But this coordination need not occur at a single place.

Coordination does need to occur at a single time, however. Every coordinated representation of the present moment constitutes a separate conscious sensory state. If there were two such representations, there would be two cases of sensory consciousness. While split brain patients may be an example of two conscious states in one brain, there is reason to believe a normal brain can only maintain one. The first consideration is the unity of consciousness. However disjointed and gappy our sensory processing is, we represent the world as unified and serially presented. A coordinated representation of the present moment accounts for this apparent unity without postulating a unity in the vehicles of representation. The second consideration is the role of sensory consciousness in decision and action. As I will argue in the next section, coordinated representations of the present moment help a creature assess current conditions and so act appropriately. Concurrent representations of the present moment would undercut this function by providing several different, and perhaps conflicting, assessments. In both of these cases, postulating multiple conscious states requires iterated levels of coordination to account for the features that simply fall out of the account if there is just one conscious state at a time.

3.1.c Decision and action

Now that we have addressed some of the concerns attendant to the idea that representations *can* 'come together', we need to ask *why* representations *would* come together into a coordinated sensory representation of the present moment. One way to motivate the distinction between representations *of now* and those that simply occur now is to consider why it might be necessary to represent the present moment. Coordination and integration are common features in theories of consciousness. Neuropsychologist Bernard Baars has proposed that conscious states coordinate the specialized work done by multiple, inde-

pendently functioning sub-processors so as to initiate coherent action (Baars 1993, 1997; Baars & Fehling 1992: 204). Robert Van Gulick also opts for an integrationist view, suggesting that the informational structure of conscious sensory states likely requires "the simultaneous interaction of many brain regions and representational systems" (Van Gulick 1992: 229). My own view highlights the role of decision and action in the formation of representations of the world at the present moment.

I have argued that the coordination of representations into conscious sensory states is geared toward accomplishing a particular sort of task, general to specific. The coordination of some sensory representations rather than others into conscious states is based on what combination will most likely help the creature in its mission. To get an idea about what factors might determine such a combination, consider a comparison of the sorts of stimuli that tend to be processed consciously with those that tend to be processed unconsciously. Baars notes that unconscious sensory processing occurs when stimuli are of short duration or low intensity, when the range of contents is limited, and when stimuli and response are habituated and routine.[96] Applied to the example above, my feeling of body cramps as I work remains unconscious because it is not terribly intense or complex, and my responses are routine. Conscious sensory processing, on the other hand, is required for long-lasting or high-intensity stimuli, wide ranges of content, or novel stimulus-response patterns. So when I cease working, my feeling of cramps becomes conscious due to a change in working conditions. I am no longer concentrating on a specific task, so other general-purpose stimuli can be incorporated into my current conscious state. The cramp stimuli become part of the complex range of stimuli that flood the mind following the release of focused attention. Additionally the need to choose a new task might generate a search for new stimuli, a search which soon hits upon the signals from leg and shoulder. Alternatively, a very intense cramp might have interrupted my computer task, calling for an immediate change in task, such as walking about a bit or stretching the offending muscle.

A common element here is that stimuli generating conscious sensory processing present an immediate challenge. They alert the creature to a potential obstacle or item of interest. Long-lasting or high-intensity stimuli suggest a dramatic change in the environment that might call for reassessment of current plans. Novel stimuli might also present motivation for altering the present course of action. Conscious evaluation of wide ranges of content may be needed to develop an overall assessment of current conditions. In my view,

the need to process and react appropriately to wide ranges of content is the main reason for the development of sensory consciousness.

As Bruce Bridgeman notes, the ability to decide between two alternative actions is the point at which a creature needs an internal system for making and executing a plan of action.

> An organism that merely reacts to sensory information has no need for consciousness – it simply does what the environment demands and its psychology is one giant transfer function. As soon as more than one plan drives behavior, however, there must be an internal rather than external trigger for action. Along with this must come a planning mechanism that makes plans, stores them in memory and keeps track of where the execution of each one stands.
>
> (Bridgeman 1992: 207–208)

The function of conscious sensory states, I maintain, is as part of the planning mechanism Bridgeman describes. Conscious sensory states are a representation of the present moment, formed in order to keep track of the relation between the organism's actions and environmental conditions. Conscious sensory states represent current conditions in order to make sure there is no new danger or opportunity on the horizon. If there is, then a change in plans may be in order to instigate new actions like avoiding that puddle up ahead or buying that delicious donut in the window. If a sensory state does not carry information required for this sort of plan assessment, then it need not be incorporated into the representation of the present moment. Only those sensory states most likely to figure in the ongoing decision-making procedure of the creature are selected by a second sense.

Antti Revonsuo has developed a similar account of the function of sensory consciousness. In describing the difference between unconscious and conscious action, he writes:

> Even apparently complex behaviors that are carried out without any conscious model of the world, such as sleepwalking and epileptic automatisms, are far from the global, flexible types of conscious control. Both of these curious instances of presumably unconscious control of behavior have been described as consisting of aimless and confused wandering, seemingly purposeful but repetitive or stereotyped, habitual patterns of behavior. (Revonsuo 1997: 189f)

Complex behaviors need to be conscious, on this view, to be appropriately 'global' and 'flexible', that is, to incorporate a wide range of information and facilitate a wide range of responses.

Thus, innate and habituated sensory-response patterns usually remain unconscious, despite the fact that they are instrumental in overall response effec-

tiveness. If there is a reason to direct attention to these otherwise unconscious sensory states, then they can become conscious. But so long as response follows automatically from sensation, there is no need to incorporate the sensation into one's representation of the world at the present moment. As noted above, sensory states all *occur* in the present moment, but they may or may not *represent* the present moment. Though sensory states represent features currently present, those features need not be included in a representation *of* the present. Habitual stimulus-response patterns, for example, need not be included in a representation *of* the present because they are not required for planning future action. When the stimulus occurs, it generates an appropriate response automatically without the need for conscious processing. One reason very strong stimuli are processed consciously, in addition to an automatic response, is their likely importance toward future plans. Though touching something very hot results in immediate withdrawal, the pain is nonetheless conscious. Because pains of this sort are quite important to planning one's next action (getting a hot pad, moving further away from the stove), the sensation has a broader impact than other simple stimulus-response connections.

Because conscious processing takes time, its value probably lies in long-term direction, rather than immediate action or even the inhibition of action. As anyone knows who has ever tried to stop herself in the middle of making a social gaffe, once an action is begun, it is nearly impossible to restrain. When lucky, there is just time enough to plan an appropriate apology. We keep track of the present moment largely in order to plan for the ones that follow. For immediate action we tend to rely on our previous plans, those informed by previous conscious sensory states and those hard-wired by evolution. Lest this remark sounds unacceptably determininist, we should not forget that we are speaking on the time-scale of seconds here. Although my conscious sensory representation of 'now' may not direct my actions at the same moment represented by 'now' – not only has 'now' passed, but actions occurring 'now' were initiated even earlier – the representation of 'now' influences my actions in the next second. Such an influence can be considerable, and it usually is.

Remember Betsy and her thimble-seeking. Betsy's ability to narrow the focus of her search was key to her success in finding the thimble. After completing the general task of determining the overall layout of the room, Betsy shifted to the more specific task of locating the thimble. This more narrow search eventually produced a sensory representation of the thimble sufficiently coordinated into Betsy's representation of the world at the present moment. That is, Betsy acquired a conscious sensory state representing the thimble. Did this conscious sensory state occur before or after Betsy began to sit down? We

may never know. The point is that sensory representations were coordinated into Betsy's conscious sensory states in order to keep track of what was happening in the world as she conducted a particular task. The constant updating of information helped Betsy adjust her responses until she eventually found the thimble.

Returning to an earlier question about how many sensory representations are required for sensory consciousness, the intervening considerations provide us with another reason to think multiple representations are involved. As argued above, creatures that are unable to make decisions probably do *not* have conscious sensory states; they simply react to sensory information according to set patterns. It is when several different actions become possible that conscious sensory states are required to keep track of current environmental conditions. Conscious sensory states inform the creature about how the world is 'now', what effects previous actions have had and whether any new prospects or dangers are afoot. So, if the function of sensory consciousness is to assess current conditions in order to facilitate effective decision and action, the number of representations coordinated into a conscious state will be determined by the interests and capacities of the creature rather than any particular logical considerations. This conclusion is supported by differential modulations in visual processing as reported by neuropsychologist John Maunsell.

> By filtering out irrelevant signals and adding information about objects whose presence is remembered or inferred, the cortex creates an edited representation of the visual world that is dynamically modified to suit the immediate goals of the viewer. (Maunsell 1995:768)

How many representations are coordinated into a conscious sensory representation? As many as are needed to accomplish the task at hand.

In conclusion let me restate the claims made thus far. The distinction between Type 1 and Type 2 cases of conscious sensory states is made in terms of degree of coordination: Type 1 states are well-coordinated sensory representations aimed toward a specific task; Type 2 states are loosely coordinated sensory representations aimed toward more general tasks. Next, the distinction between Types 1 and 2 (conscious) and Type 3 (unconscious) sensory states is made in terms of their relation to a coordinated representation of the present moment. Conscious sensory states are so coordinated, and unconscious sensory states are not included in these coordinated representations. The proposal is that a conscious sensory state is constituted by the following necessary and jointly sufficient conditions: (1) a coordinated representation of the world,

(2) which includes a representation of the present moment, (3) produced by a second sense.

3.2 What good is a second sense?

If you found Part 1 as persuasive as I hope, you may be wondering why I call this the *second sense* theory instead of something like the *present tense* theory. The critical element in distinguishing conscious from unconscious representations is that the conscious ones include a representation of the present moment. What does a second sense add to the theory?

We still have not answered one of the critical questions in determining the difference between conscious and unconscious sensory states: how is it that some sensory states are conscious and others are unconscious? We now have an idea about what constitutes the difference between conscious and unconscious sensory states and about why there exists such a difference. What remains is to show how some sensory states come to be conscious, for if we fail to address this additional problem, then we have failed to explain sensory consciousness and have merely analyzed it. An adequate analysis would be no small achievement in itself, given the disagreements and inherent murkiness of the topic. Yet an explanation of how sensory consciousness comes about is necessary to dispel the sense that it is a mysterious entity, somehow beyond our full comprehension. My efforts up to this point have been to identify a phenomenon to be explained, sensory consciousness, and to offer a theoretical definition of that phenomenon. The next step is to provide an account of how conscious sensory states come about.

I suggest that a second sense selects and coordinates sensory states to form the coordinated representations of the present moment that are conscious states. Borrowing an insight from higher-order inner sense theory, I maintain that some kind of sensory mechanism is responsible for making sensory states conscious. The following offers a causal explanation, yet it is still part of the project of describing in what conscious sensory states *consist*. On my view, conscious sensory states are coordinated representations of the present moment, and to be so coordinated requires the operation of a second sense. A second sense selects sensory representations and combines them to produce coordinated sensory representations of the world at the present moment, which are conscious sensory states.

A comparison of this view with the higher-order inner sense theory may help illuminate the role of the second sense. According to the higher-order the-

ory, an inner sense produces coordinated representations *about sensory states*. Sensory states become conscious by virtue of the higher-order representations about them, which are produced by an inner sense. What it is to be a conscious state is to be represented by an appropriate higher-order representation, and what it is to be an appropriate higher-order representation is centrally to be the product of an inner sense. According to the second sense theory, the second sense produces coordinated representations *about the world* at the present moment. Sensory states become conscious by being coordinated into a representation of the world at the present moment, produced by a second sense. What it is to be a conscious state is to represent the world at the present moment, and what it is to represent the world at the present moment is centrally to be the product of a second sense.

A second sense is necessary for producing conscious states in order to account for the particular form of coordination involved in sensory consciousness. In distinguishing between 'well-coordinated' and 'loosely coordinated' representations, I argued that the difference is task-related. Specific tasks require well-coordinated representations, while general tasks can be conducted with only loosely coordinated representations. Representations are coordinated in accordance with the task at hand. The second sense is the mechanism required to select and combine the representations needed to perform the task. Sensory states are conscious when selected by the second sense and coordinated into a representation of the present moment. Sensory states not selected remain unconscious.

One often overlooked feature of conscious sensory states is that it is possible to control their content to some extent. In exercising this control, sensory consciousness is intimately connected to attention.[97] By a shift of attention some things rather than others become the content of our conscious sensory states. Though I am not normally state conscious of my body position, I can become state conscious of how I am sitting, the tilt of my head, the seat pressing against my legs, etc. by focusing on these body positions. Similarly, I can shift my attention from the computer screen to the hum coming from the next room, thereby becoming state conscious of the refrigerator sound. Thus, previously unconscious sensory representations – of body position, of background sounds – become conscious through purposive shifts of attention. Because sensory consciousness (as well as attention) is a matter of degree, there may be no way to determine whether a particular attention shift has made an unconscious sensory representation conscious or whether a loosely coordinated conscious sensory representation has merely become well-coordinated. That is, we may not be able to say definitively whether a Type 1 case was formerly a Type 3 or

Type 2. Whatever the case, the shift is the result of the functioning of a second sense. The ability to direct attention so as to be state conscious of some things rather than others is accomplished by a second sense.

Because conscious sensory states are coordinated sensory representations according to a particular sort of task, there must be some way to control which representations are selected. In many situations, some sorts of sensory representations are clearly irrelevant or downright distracting. If there were no way to select some sensory representations and eliminate others, sensory consciousness would be overloaded with a random mix of sensory content. Think again of Betsy and her thimble-seeking. Even before she was able to focus her search sufficiently narrowly to identify the thimble, she had many sorts of sensory representations that remained unconscious. Some examples of Betsy's unconscious sensory representations were her representations of body position and the subtle motor adjustments that allow her to avoid obstacles in the room. These representations were peripheral to her current task and so remained unconscious. Which sensory representations are conscious and which remain unconscious is task-relative, so there must be some way to select which sensory representations need to be coordinated. Without this sort of control over the content of conscious sensory states, their value in accomplishing tasks would be lost. For conscious sensory states to be effective, there must be something like a second sense to select task-appropriate sensory representations.

Of course, the content of conscious sensory states is not always under our control – a nagging tune or disturbing image are familiar examples of uninvited conscious contents. Such cases raise the question of who controls a second sense. While it is possible for us as conscious creatures to exercise a great deal of control over the contents of our conscious sensory states, it is useful to keep in mind that sensory consciousness is much more basic evolutionarily than the sophisticated forms of consciousness we have developed, such as self-consciousness or social consciousness. As Bridgeman suggests, sensory consciousness probably arrives with the ability for deciding among possible actions (Bridgeman 1992: 207–208). Oftentimes the choice of action will be based on the creature's current goals and desires, and so the content of conscious sensory states will reflect those purposes. But conscious content is not always a matter of top-down control. At other times, the urgency of the stimulus may override other creature wishes, such as when feeling severe pain or hearing a sudden, loud sound. The evolutionary advantage of this sort of bottom-up control of conscious content is fairly clear. Whatever task the creature had planned may need to be abandoned in light of this new information about the world. The value of a nagging tune or disturbing image is harder to explain. But such in-

trusions may exploit the usefulness of other sorts of representations that have more obvious survival value. Nagging tunes tend to nag because they have been repeated so often, and repetition is a useful way to sort between transitory things and things that recur. Recurring things are good to remember because they can serve as markers for both positive and negative features of the environment. Likewise, really shocking things are good to remember – the shock is usually quite pleasant or unpleasant – so it is handy to easily recognize and pursue or avoid them. Whatever the final explanation, we should consider the control function of a second sense in larger terms than the satisfaction of a particular creature's current goals and desires. The second sense determines the contents of consciousness primarily on the basis of evolutionary value, and only secondarily according to the creature's goals.

It is also worth pointing out that a second sense is in this way similar to the external senses. I can send my eyes or limbs on a specific mission to acquire information about one or another aspect of the current environment, yet I cannot fail to see what is directly in front of my eyes or feel what my limbs are touching. I may fail to see *that* it is an apple in front of me or feel *that* the substance in the bowl is macaroni, but I cannot fail to sense the stimuli that impact my sensory organs. In this way, the senses are passive. Yet they are also active and can be to some extent directed, with sight and touch as most amenable to direction, smell and taste least directable, and hearing somewhere in between. I can hear some things better by turning my head, or cupping an ear, but by and large I will hear any sounds in range. A second sense would likely fall in line with sight and touch on the continuum of control. When deeply focussed on a task, I am able to eliminate almost all distractions. Yet when first waking in the morning, I am bombarded by sensations with little ability for selection or control.

Given the analogy here between external senses and a second sense, this is a good point to consider the reasons to call this mechanism a 'sense.' In what ways is a second sense similar to the external senses, and in what ways is it different? In Chapter 1 I listed three features essential to a 'sense': (1) it is non-cognitive; (2) it serves a relay function; and (3) it has particular forms of inputs and outputs. All three features apply to a second sense.

First, the mechanism in question is *non-cognitive*. That is, sensation is non-conceptual, involves tracking, and is detailed, whereas cognition is conceptual, represents in absence and abstracts from detailed presentations.[98] The second sense exhibits all three of these markers for sensation, but differs from external sensation in one important respect. For external senses, the causal source of sensory input is the same as the object represented in sensation. If I have

a sensory representation of an apple, provided the representation is veridical, the causal source of that representation is an apple. I track the apple with my eyes, and the detail represented is about the apple. For the second sense, however, causal source and object represented come apart. Sensory states are the immediate causal source of conscious sensory representations, but they are not the object represented, on this account. So if I have a conscious sensory representation of an apple, the causal source of that representation is a sensory representation of the apple.

This difference may not seem terribly significant, since there are causal intermediaries in all cases of sensation. My sensory representation of the apple is not immediately caused by the apple, but by light waves, retinal activation, neuronal stimulation and a host of other intermediary causes. However, these intermediary causes are not themselves mental representations. In second sensing, sensory representations are among the causal intermediaries between the apple and my conscious representation of the apple. Furthermore, the *content* of sensory representations is tracked by a second sense in order to determine its appropriateness to the task currently underway. This point raises the question of how a second sense can select sensory representations for coordination into a conscious state without doing so by means of a higher-order representation. I will come back to this question in a moment when I consider some of the technical difficulties in explaining how a second sense produces conscious sensory states.

The second reason to consider a second sense to be a 'sense' is that it functions primarily as a relay mechanism. External senses take various forms of physical stimuli as input and relay this information in the form of sensory representations to cognitive structures, motor systems, and, per hypothesis, to the second sense. Similarly, the second sense takes sensory representations as input and relays this information in the form of conscious sensory representations to cognitive structures, motor systems and in all likelihood, back to the external senses. Being a relay mechanism alone is clearly not sufficient to qualify anything as a 'sense' since cognitive structures rely information and are by definition non-sensory. Nonetheless, the relay function is a necessary feature of a second sense, and in conjunction with the other two features of a sense serves to distinguish it as a sense.

Third, the flip side of the previous point, a second sense has the function of relaying specific inputs by virtue of a specific form of output. External senses have the function of relaying various forms of physical inputs, such as light waves or sound frequency. The function of a second sense is to relay sensory representations. The output of each sense is also specific, producing its own

variety of representation. Eyes produce visual representations, ears produce auditory representations, and so forth. Similarly, the second sense produces conscious sensory representations, coordinated sensory representations of the world at the present moment. Conscious sensory representations combine features from several sensory modalities and so are not unique in the *kind* of feature they represent. However, they are unique in the *way* they represent those features, as coordinated sensory representations of the world at the present moment. No other representations represent in just this way.

There are no doubt other differences between external senses and the second sense,[99] but the similarities are sufficient at least to say that the operations involved are more like sensation than cognition. The ability to control the content of conscious sensory states suggests that there is a mechanism of some sort dedicated to producing them. This mechanism does not require concepts, relays a specific sort of information (sensory representation), and produces a specific sort of representation (conscious sensory representation). For these reasons, it seems justifiable to call the second sense a 'sense'.

Three questions remain about exactly how a second sense produces conscious sensory states. The first question concerns the mereological relation between unconscious and conscious sensory states. I see two possible combinations, either of which might be the case. On one hand, it could be that the very same states that were unconscious compose a representation of the present moment. In this scenario, an unconscious sensory state becomes conscious when selected by a second sense and coordinated with other formerly unconscious states to form a new, composite state. The content of the new state includes the contents of the states that compose it, and by virtue of the selection and coordination process, it acquires the additional content of 'now'. Alternatively, a second sense could produce an entirely new complex representation with the same content as the unconscious states, but not composed of them. In this case as well, the complex representation would have the additional content of 'now.' Earlier I noted my preference for the compositional relationship as it seems more economical, but the thesis is consistent with the second alternative as well. The key element in both descriptions is that the composite or complex state represents the present moment by virtue of the selection and coordination process of its contents. So we see that the second sense is not merely peripheral, but produces a representation of the present moment through its selection and coordination function.

The second question, raised above, is how a second sense can select and coordinate sensory representations without higher-order representations. Given that the second sense coordinates sensory representations from several sources,

there must be some way for it to identify the content of the representations in order to get the combinations right. Suppose I am currently representing a variety of colors, sounds and shapes, all of which get coordinated into my conscious sensory representation of a train passing. This happens to be a well-coordinated representation because I am running late and so have been focusing my attention on the tracks ahead. The question is, how does the second sense select from all of the available sensory representations to produce this useful, albeit disappointing, conscious sensory representation of the train passing? How does the second sense determine which are the sensory representations of the train at this moment? To do this, the second sense must identify what the sensory representations are about and when the events represented occurred relative to other events represented by other representations. But identifying a sensory representation as one about a train at a moment seems to require representing the state as having this content. Thus, it would seem that the second sense must produce higher-order representations, representations about its sensory representations, in order to accomplish its coordination function. If so, the second sense cannot be flat after all; it must be higher-order.

While representing one's representations is certainly one way to coordinate them, it does not seem to be the only way. Representations could be coordinated by means of their vehicles. That is, it might be the case that features of the content of representations are somehow encoded in their vehicles, and the second sense is attuned to these vehicular properties. If so, nature offers multiple models for how selection and combination might be performed non-representationally. In photosynthesis, for example, plant cells extract water, carbon dioxide and solar energy from the atmosphere to produce complex hydrocarbons for the plant's growth and development. Or consider a more sophisticated selection and combination mechanism, the computer. Computers are selection and combination wizards but any sensitivity to content is based purely on its syntactic structure. To the extent that my computer represents 'consciousness,' it does so by some string of 0s and 1s. If the computer searches for 'consciousness', it will do so by selecting instances of the vehicle for representing 'consciousness' in the computer, the string.[100] I see no reason why the second sense could not be attuned to non-representational features of sensory states in a similar way. The difficult part of the story is to give a satisfying account of how representational content might come to be encoded in the vehicle. For the vehicle must wear its content on its sleeve in order to be sorted by the kind of 'stupid' mechanism I propose. I can offer no specific proposal about the sort of coding system that might work here, but there seems no principled objection to such a system.[101]

Moreover, even if something like a higher-order representation were required to coordinate sensory representations, the resulting theory would still differ significantly from existing higher-order theories. For the sake of argument, say a second sense does indeed require higher-order representations in order to perform its coordination function. Note that the conscious sensory states which are the *output* of this new hybrid sense still represent the world; they are the result of coordinating sensory representations about the world. On this hybrid sense theory, as in my second sense theory, conscious sensory states are coordinated sensory representations of the world at the present moment. The hybrid sense utilizes higher-order representations in order to produce flat representations about the world. By contrast, traditional higher-order theories claim that a sensory state becomes conscious when represented by a higher-order state; the representation relation itself constitutes sensory consciousness. There is no new coordinated representation of the world, according to higher-order theorists, only the original sensory representations now enjoying the dubious honor of being represented.

The hybrid theory seems to be an unwieldy combination of flat and higher-order theories and so I will not consider it further. I suggest it only to show that, even if we admit the need for higher-order representation in producing coordinated sensory representations, this is not in itself an argument for either higher-order thought or higher-order perception theories as they currently stand. Instead of exploring a hybrid theory, we should think more about how a truly flat second sense might work. It is an engineering problem: given these sorts of input, how could a device produce these sorts of output? I wish I had a ready answer, but I am willing to bet that more technical minds than mine could find one.

The third question addresses a lingering doubt that often haunts attempts to offer a substantive explanation of consciousness: isn't the second sense some kind of homunculus? Doesn't the selection process of the second sense require that there be a 'little person' who 'sees' the representations and thereby chooses which ones to collect into a conscious sensory state? This sort of homunculus would indeed be problematic, for then we would need to explain how our internal homunculus 'sees' the representations and chooses them. We would then have the same problem of sensory consciousness, once removed. As Dennett notes, the key to avoiding the problematic sort of homunculus, call it the *Cartesian homunculus*, is to insure the mechanisms involved are "relatively stupid functionaries" (Dennett 1991:14). As I suggested in response to the previous question, the second sense may operate in a way similar to the 'relatively stupid' functional operations of photosynthesis and computer pro-

cessing. There are no infinite regress worries concerning these sorts of selection/combination mechanisms, so to the extent that the second sense runs on similar principles, it is immune from such worries as well.

Now, if we cannot find a purely flat engineering solution to the selection question, the second sense will necessarily be a much smarter sort of homunculus. On the hybrid view, the second sense must produce higher-order representations of sensory representations in order to determine which representations to select. Yet even here the mechanism can be sufficiently stupid to avoid the regress of the Cartesian homunculus. Lycan conceives of a higher-order sensor as "an internal scanner or monitor that outputs second-order representations of first-order psychological states" (Lycan 1996:31). There need not be an 'executive' scanner, and scanners can be directed at "representational subsystems and stages of same" (Lycan 1996:32). The result is a decentralized model of organization that Dennett has characterized as 'pandemonium,' where specialist mechanisms compete for control of mental processes (Dennett 1991:239). Lycan approvingly describes Dennett's "'Joycean' machines that formulate synthesized reports of our own passing states" (Lycan 1996:31). Though I believe the content of conscious sensory states is more coherent than the pandemonium model suggests,[102] it is clear that neither the flat second sense nor the hybrid sense requires a Cartesian homunculus to do its job.

There you have the complete account of the second sense theory of sensory consciousness. So let us review the main elements of the theory. I have argued that conscious sensory states are coordinated sensory representations of the world at the present moment. A second sense is the mechanism that selects sensory representations and coordinates them into representations of the present moment. Therefore, the operations of the second sense are necessary to determine which sensory states become conscious and which do not. This selection function of a second sense explains the way a creature can (to some extent) control the content of her conscious sensory states. Though sensing is not like external sensing in every respect, there are sufficient similarities to call this mechanism a kind of sense. It is non-cognitive, relays information, and takes a particular form of input and output.

Though the theory introduces no new elements, its combination of elements and emphasis on certain features such as the temporal dimension of sensory consciousness is original. As such, there are no doubt scores of objections and refinements to the theory that lie ahead. I have considered a few and will review some more general objections in Chapter 5. The next section will revisit the issue of Spot-sight and thimble-seeking to see how the second sense theory deals with this problem.

3.3 Spot-sight again

With the second sense theory in hand, let us return to the problem of Spot-sight and thimble-seeking. Dretske's Spot, the unnoticed difference between object sets Alpha and Beta, and the unnoticed thimble that sat directly in front of Betsy's nose are cases of a troublesome category for explanations of sensory consciousness. They are cases where one seems to be state conscious of an object without noticing it. As listed earlier, the three basic types of sensory states that a theory of sensory consciousness needs to explain are:

1. being state conscious of *x* and noticing/responding to *x*
2. being state conscious of *x* without noticing *x*
3. not being state conscious of *x* but nonetheless able to respond to *x*

Standard cases of sensory consciousness (Type 1 cases) involve noticing the object of one's conscious sensory state, whereas unconscious sensory states (Type 3) involve no such notice. Spot-sight and thimble-seeking (Type 2 cases) do not seem to fall neatly in either category, rather they are somewhere in-between. Characterizing this 'between-ness' has been the project of this chapter.

I have argued that Spot-sight and thimble-seeking are cases of conscious sensory states which are more loosely coordinated than cases where one does notice specific objects. The difference between Type 1 and Type 2 cases, in my view, is one of degree of coordination. In Type 1 cases, sensory representations are well-coordinated in order toward a specific task. One is able to select and integrate sensory representations about particular objects so as to respond appropriately. Type 2 cases involve more loosely coordinated sensory representations for more general tasks. Such loosely coordinated sensory representations are useful in assessing broad structural features of the environment or when no specific task is at hand. More than degree of coordination is needed to account for the distinction between conscious sensory states (Type 1 and 2) and unconscious sensory states (Type 3), however, as unconscious sensory states often exhibit a high degree of coordination. I have therefore argued that we think of conscious sensory states as coordinated sensory representations of the world at the present moment. While unconscious sensory states may be coordinated sensory representations of the world, they are not included in one's representation of the present moment. Unconscious sensory representations do not represent the world as 'now.' Finally, a second sense is the coordinating mechanism that selects which sensory states are incorporated into a representation of the present moment. If a specific task is at hand, sensory representations related to that task will be coordinated into a Type 1 state. If more general

tasks are underway, a less selective process will apply. Sensory representations not selected for coordination by a second sense may be coordinated in some way, but they will not be coordinated into the representation of the world at the present moment.

Now let's apply this theory of sensory consciousness to the two trouble cases from Chapter 2. I have already considered the main features of Betsy's case, but it may be helpful to go through them all at once. As Betsy searches in vain for the thimble, her conscious sensory states are largely of the Type 2 variety. Though she has a specific task in mind, spotting the thimble, she must first undertake a more general survey of her environment. Only once she has navigated around the room a few times and is sufficiently familiar with its structural features will she be able to narrow her focus to thimble-sized objects. At this point she undergoes several Type 1 states as she notices and rejects several thimble-candidates, until she finally notices the target object and sits down. Throughout this adventure, Betsy will undoubtedly have many unconscious sensory representations, both coordinated and uncoordinated. Some of her more coordinated unconscious representations include her perceptual assessment of distances to various obstacles (so as not to bump into things) and proprioceptive feedback (to regulate her walking and head movements and such). Betsy's second sense is charged with the arduous selection process required to balance coordination requirements with task requirements in order to present Betsy with coordinated sensory representations of the world at each passing moment of play.

Regarding Spot-sight, the story is only slightly different. When examining Alpha and then Beta, we begin with Type 1 states. We are told to look at every part of Alpha and Beta, and the peculiar nature of the request cues us to pay particular attention. Well-coordinated sensory representations of Alpha and Beta result. But then Dretske points out Spot. While Spot is part of Alpha, chances are that few viewers noticed Spot in particular. So Dretske's call to notice Spot necessitates a yet more well-coordinated sensory representation. The more specific task of picking out Spot requires that the coordination processes of the second sense be even more selective.

In this case only visual representations are coordinated, so it may seem that the work of Spot-sight could be done by vision alone, without a second sense. But it is important to remember that vision isn't a single, straight-shot process from retina to cortex. It involves at least two pathways and four forms of representation (Zeki 1993). So even if only visual representations are involved, coordination between, say, shape and color, is still required. Also, isolating Spot requires eliminating distractor information. The second sense both combines

and *selects* sensory representations for inclusion into conscious sensory states. Without the selection of some visual representations over others, there would be no way to locate Spot amidst the other spots and shapes. So when we go back to the figure, sure enough, this time we spot Spot. Now we can identify Spot as Spot rather than take Spot as part of Alpha. Note that the ability to identify Spot *as* Spot is facilitated by the coordination processes of sensory consciousness but is not required for sensory consciousness of Spot. On the present theory, we were state conscious of Spot throughout. Focussing more narrowly on Spot allowed Spot to be identified, but it was represented by a conscious sensory state all along. The sensory representation of Spot as part of Alpha was part of the initial conscious sensory state. By including the representation of every part of Alpha in the coordinated sensory representation of the world at that moment, Spot, as part of Alpha, was included as well. We were state conscious of Spot, yet we failed to notice Spot.

One can continue to deny that Type 2 cases exist, and staunchly maintain that we are state conscious of no more than we notice. But this position seems at odds with the intuitive sense we have that our conscious sensory states, like our sensory states, are more richly detailed than we are able to fully process. We might simply represent there to be more detail without wasting *paint* on details, as Dennett and Kinsbourne put it.[103] But this answer fails to satisfy. The only reason I can see for withholding the admission that we are state conscious of more than we notice is the worry that we could not account for the cases of sensory consciousness without noticing. Since I have provided an account of these troublesome cases, I suggest we embrace them. Let a thousand loosely coordinated sensory representations bloom.

All is not settled for the second sense theory, however. A final difficult theoretical problem remains to be discussed, the problem of subjectivity. Sensory consciousness requires a subject of consciousness; there must be someone for whom the world is represented as appearing at the present moment. Without an account of the subject of consciousness, homunculus concerns such as those I dismissed above could return in full force. Someone could concede that conscious sensory states are as I have described but argue that the *real* problem of consciousness lies with the subject of consciousness. As I argued in Chapter 2, exactly this ambiguity about the location of the mystery of consciousness compromises higher-order theories of consciousness. To block a similar charge against the second sense theory, I need to give a completely non-mysterious account of subjectivity. My goal is to show that, at root, subjectivity is quite bare; indeed it is antecedent to sensory consciousness. Therefore, whatever

mystery arises with sensory consciousness cannot be shifted to the subject of consciousness.

There is still some bumpy road ahead, but the main obstacles have been passed. In view of the limitations of higher-order theories of state consciousness and Dretske's flat theory, I have proposed a new second sense theory. On the proposed theory, conscious sensory states are representations of the world at the present moment, selected and coordinated by a second sense. Conscious sensory states may be well-coordinated, as when one is focused on a specific task, or loosely coordinated as when a more general task or no particular task is underway. Unconscious sensory states, though they may be coordinated, are not coordinated into one's representation of the present moment. Thus, the theory accounts for all three basic types of mental states that a theory of state consciousness needs to explain: being state conscious of *x* and noticing/responding to *x*, being state conscious of *x* without noticing *x*, and not being state conscious of *x* but nonetheless able to respond to *x*.

Chapter 4

Subjectivity

One of the reasons conscious sensory states have seemed so mysterious to some is that they are subjective, experienced somehow from the 'inside'. Only the subject of a conscious sensory state experiences that state, and that particular sort of experience constitutes a phenomenon that seems wholly unlike any other sort of phenomenon. When unconscious, the sensory processes from stimulus to response submit to familiar sorts of mechanical and electrical explanations. There seems to be nothing otherworldly about the way a tap on the knee causes the foot to jump. Now add to this banal scenario the fact that it is *my* knee and *my* foot,[104] and that I experience the tap and the jump from a *point of view* that no one shares. What is this point of view that only I possess? How is my point of view related to the fact that it is my knee and my foot? These are questions about subjectivity, and they have been particularly nettlesome in philosophy of mind, because it seems as though something like a point of view cannot be accommodated by scientific theory. Science undertakes objective study; it requires objects that can be examined by multiple observers in a variety of situations. Subjectivity is by definition not objective and so cannot be explained scientifically. Thus, subjectivity is a phenomenon outside the realm of scientific study. Thomas Nagel classically stated the problem this way: "every subjective phenomenon is essentially connected with a single point of view, and it seems inevitable that an objective, physical theory will abandon that point of view" (Nagel 1974/1991:423).

As an epistemological point, the argument stands. Things known from a particular subject's point of view, such as her conscious sensory states, are known in a way unavailable to anyone other than the subject. More will be said on the nature of this privileged access in a moment, but note that as a purely epistemological point it is in no way a challenge to materialism. What often happens to make the argument anti-materialist is that the notion of subjectivity in the conclusion gets reified into a metaphysical entity, and this entity purportedly cannot be explained in terms of physical theory. One project of this chapter will be to tease apart the epistemological claims about subjectivity from the metaphysical ones in order to show that the completeness of physicalist ex-

planations of mentality is not threatened by subjectivity. But before we can get to this task, we will need to clarify the notion of subjectivity at issue. Like many issues in the contested area of consciousness research, developing a working definition of the phenomenon in question is the first order of business. What I take to be subjectivity proper is, roughly, having a point of view, a unique perspective on the world, not assumable by anyone else. Though this description requires refinement, it should help us begin to distinguish subjectivity from related concepts such as: qualitative character,[105] first person authority and privileged access.

Undoubtedly some will object that the very idea of isolating subjectivity from the net of related concepts is flawed. Rather than providing clarity, the result will necessarily be a deflated notion of subjectivity empty of the key features that make subjectivity a problem in the first place. Without a full-blooded notion of subjectivity, the objection goes, any explanation will be trivial or at least uninteresting. To such as these, anything unmysterious is uninteresting, and since my aim is to dispel the mystery of subjectivity, a successful result will no doubt disappoint. The explanation will not be trivial, however, as it will provide an account of subjectivity that is unmysterious, compatible with physicalism, and evolutionarily plausible. The cumulative effect of a hodge-podge of notions conflated together makes subjectivity seem incompatible with physicalism. So by applying a bit of language therapy, we can solve the problem we ourselves have created. What seemed like a problem is largely a confusion of language.[106] By rolling together the problems of qualitative character, first person authority and privileged access with subjectivity, we have made all of them intractable.

That said, I should also say that my project is not an eliminativist one. Though subjectivity should be distinguished from qualitative character, first person authority, and so forth, I believe these other phenomena require a separate account. Moreover, I take it that subjectivity itself raises substantive questions that require an account. The method is nothing more than the divide and conquer strategy used earlier to consider the various phenomena under the term 'consciousness'. My strategy is to carefully distinguish each element conflated into a term in order to more easily explain them. Analysis need not result in eliminativism, unless we happen to discover that nothing was there in the first place.

The first distinction involves two of the most slippery concepts: subjectivity and qualitative character. The collapse of these two concepts appears in phrases like 'the subjective character of experience,' where the topic is how object *x* looks to subject *y*. But in the course of discussion, subject *y* invariably

drops out and the entire argument concerns how object *x* looks, its qualitative character. What makes green look different from red? Could what looks red to me be what looks green to you? Why does it seem impossible to describe how things look to me? All are questions about how things *look*. We assume there has to be a subject for there to be qualitative character, so subjectivity and qualitative character are clearly related. But the nature of the subject, what makes her a *subject*, is rarely considered in the discussions of how things look. In other words, the *subjective* aspect of qualitative character does not come up.[107] As we will see in Part 2, materialist arguments about qualitative character generally provide reasons to believe that the physical nature of qualitative character can be known both subjectively and objectively, and there the arguments stop. Having made qualitative character safe for science, many conclude that the problem of subjectivity is rendered moot. It is not. Even if we believe, as I do, that qualitative character admits of objective description,[108] we are still a long way from understanding how or why qualitative character has a subjective aspect or how subjective description differs from objective description.

Part 3 will attempt to advance our understanding of subjectivity by looking specifically at the distinction between subjective and objective as provided by inner sense and second sense theories. The first account by Lycan utilizes resources from his higher-order inner sense theory to argue that subjectivity is an exclusive language of thought composed of higher-order representations. While the Lycan account is characteristically intriguing, I suggest that, like the higher-order account of conscious sensory states, it is too sophisticated to account for the most basic level of subjectivity. As a foundation for the higher-order theory, I will present a Gibson-style account. Following Gibson, I assume a fundamental link between subjective and objective aspects of sensory information in that objective information is mapped relative to a subjective coordinate system. Following up, Part 4 will use the Gibsonian theory to illuminate the related problem of ownership of one's sensory states.

4.1 Subjective authority

One often hears that there is no distinction between appearance and reality for the subjective point of view. From the subjective perspective, the way things look is the way they are. Appearance is reality from the subjective point of view because subjectively one only has access to the appearance of things and appearances appear only to a subject. Hence, if appearance is reality, one cannot be mistaken about appearances. The subjective perspective yields a special, in-

corrigible knowledge about the way things look. Or so the story goes. Incorrigibility claims such as this one have taken many forms and have been subject to at least as many refutations. But the force of the claim is so strong that it continues to return, phoenix-like, in arguments about the epistemic authority of the subjective perspective. Even though I believe the incorrigibility claim is unwarranted, its force calls for explanation or the claim will simply reappear in another guise at the next opportunity.

One source of the power behind the argument for incorrigibility lies in the grammatical structure of first-person statements. 'Subjective' and 'first person' are often used interchangeably, resulting in mistaken claims about the immunity from error of subjective reports. Confusion begins with the grammatical fact that in direct discourse first-person terms refer to the grammatical subject. Whoever uses a first-person term this way cannot fail to refer to herself. My use of 'I' successfully picks out me and only me, and only I can use first person terms to pick out me. Add to this the fact that I use first person terms when I describe observations from the subjective perspective – I see, I feel, I think – and it is a short step to the idea that descriptions using first person terms cannot fail to describe me and my states. Just by using 'I', the grammatical necessity of direct discourse ensures that I will be successful in attaching my description to my self. Indeed, this much is true.

The mistake enters when we confuse the grammatical subject with the epistemic subject. It is a matter of grammatical necessity that I will refer to myself when I use first person terms in direct discourse, but this fact carries absolutely no epistemic weight on its own.[109] Grammatical necessity does nothing to ensure the accuracy of descriptions I assign to myself, nor does it ensure that I have any knowledge whatsoever about the referent of first person terms. Just as I can use the word 'dog' to successfully pick out dogs without knowing anything at all about dogs, I can use the word 'I' to pick out myself without knowing anything at all about myself.[110]

The self-referential aspect of first person terms is a fact about the language, not about me and my epistemic capacities. My ability to use first person terms to describe myself and my states is certainly convenient, but endows first person statements with no special epistemic status whatsoever. If first person statements have any special authority it must be rooted in some other source than the grammatical self-reference of first person terminology. In an attempt to avoid the deep confusion between 'subjective' and 'first person', I will restrict myself to terms like 'subjective authority' and 'subjective point of view' and refrain from the use of 'first person' in this context.

Putting grammar aside then, there is still the sense that descriptions from the subjective perspective carry a special authority not granted to objective descriptions. When I describe how things appear to me, it seems that I cannot be mistaken in describing the object's appearance. While I may be wrong about the object's *real* features, I cannot be wrong about its *apparent* features. From the subjective point of view, the distinction between appearance and reality seems to collapse. Again, there is a grain of truth to this line of argument, but upon analysis the seed falls fallow.

Immediately when presented with an argument about appearances we should ask ourselves *what* exactly an appearance is supposed to be. If the answer is that an appearance is an object and/or its properties which are present before the mind, appearances soon take the form of apparitions. By granting appearance any metaphysical substance, as things that they appear to be, the problems with sense-data theory arise in all their force. Where are objects of appearance located? What are they composed of? Appearances are often indeterminate as between, say, being round and square, but how could objects be indeterminate? Frightening objects they would be. Rather than travel down this dark and unfruitful path, we can say that appearances are not metaphysical objects but *intentional* objects. An appearance is the intentional object of a representation; it is what the mind represents there to be.[111]

So to say an object is as it appears from the subjective perspective is to say that it is as the subject represents it to be. The appearance/reality distinction collapses in the sense that, from my subjective point of view, appearances are the way I take them to be. The way things are from my perspective is the way I represent things to be, and the way I represent things to be is the way they are from my perspective. Subjective authority is grounded in the fact that I am the bearer of my representational system, and I am therefore in the best position to say how I represent things to be.[112] I assert this authority, notably, when I introspect. I represent how things appear to me; I represent how I represent things to be.

Thus the claim that there is no appearance/reality distinction from the subjective point of view could be one of two claims, neither of which is true. The first claim is that there is no distinction between appearance and reality because the subjective point of view reveals the same appearance whether the appearance is veridical or not. Reality, being objective, does not figure in the subjective point of view. Here the representationalist will concede that the appearance is the same in both veridical and hallucinatory cases, since the intentional object is the same. But this fact does not render reality irrelevant to appearances. Like any sort of mental representation, appearances can be either true or false rep-

resentations. Though the appearance is as it is represented (a point that is now a tautology), it may fail to represent what it has the function of representing. After-images are a good example. Usually the appearance of a green expanse in the visual field represents a green-colored physical object or light source at a particular external location. But when having an after-image, there is nothing external for the appearance to represent. Something has gone wrong with the representational system.[113] Though I represent a green-colored expanse, this representation fails to do what it is supposed to do – represent green *things*. Furthermore, the possibility of misrepresentation exists even if the representation is about one of the subject's own bodily states. While I may be right in saying that it seems to me that my toe hurts (or, more technically, I represent my toe as damaged), I could be wrong about what bodily state I am actually in. Though I represent my toe as hurting, it may in fact be cold.[114] To accept that appearances are the intentional objects of representations, a move I heartily endorse, is to accept that there is a reality behind every appearance, a reality that may be misrepresented. My subjective authority with respect to appearances does not extend to the states of affairs they represent, even if those states of affairs occur within my own skin.

The second claim that there is no distinction between appearance and reality from the subjective point of view is a higher-order claim about introspection. The claim is that when I represent how things appear to me, I cannot be mistaken. In introspection reality does seem to drop out, since all we introspect are the appearances. Even if there isn't *really* a green spot, there *appears* to be a green spot. And if there appears to be a green spot, surely it must appear to me that there appears to be a green spot. How could the green spot appearance appear otherwise than as a green spot appearance? The reality of being a green spot appearance is just its appearing to me to be a green spot appearance. Again, the answer is yes and no. The reality of being a green spot appearance is that I represent there to be a green spot. In other words, there must appear to me to be a green spot in order for me to have a green spot appearance. But it does not follow that I represent myself as having a green spot appearance every time I have one. I can be mistaken about how things appear to me when I introspect because, on my view, introspection is just another form of representation and so can fail. Even if there appears to be a green spot, it may not appear to me that there appears to be a green spot. I may represent there to be a green spot without representing that I represent there to be a green spot. A defense of this view of introspection would take us too far afield to deliver here, so I will simply draw two limited conclusions. First, the possibility of an alternative means that the claim of no appearance/reality distinction from the subjective

point of view must be argued rather than assumed. Second, the claim applies, if at all, only to the objects of introspection. Regarding conscious sensory states, the distinction between appearance and reality has exactly the same force as it does in any case of representation.

4.2 Special facts or special access?

Even if introspection is fallible, it still seems that there is something special about my subjective perspective on my conscious states. My mental states are special because no one but me knows what it's like to have them. This claim is best known through Nagel's comparison between what it's like to have human mental states and what it's like to have bat mental states, a comparison that has received considerable attention since its introduction (Nagel 1974/1991). Consequently, most of what I have to say is not new.[115] A review of two main objections to Nagel's arguments will nonetheless clear away the final vestiges of concerns about qualitative character so that we can concentrate on the distinctive problem of subjectivity itself.

4.2.A On what it's like

Nagel's argument is important to examine because it moves clearly from observations about 'what it's like' to have qualitative character to a description of subjectivity. The move can be broken up into four steps.

1. Nagel begins with consciousness[116] and claims that qualitative character is central to conscious states: "fundamentally, an organism has conscious mental states if and only if there is something that it is like to *be* that organism" (Nagel 1974/1991:422). The 'what it's like' locution is notoriously vague, and much of the persuasiveness of anti-materialist 'what it's like' arguments can be attributed to this indeterminacy. But qualitative character is certainly a prominent, if not exhaustive, element in Nagel's notion of what it's like for the organism. Take, for example, the charge to physicalists: "If physicalism is to be defended, the phenomenological features must themselves be given a physical account. But when we examine their subjective character it seems that such a result is impossible" (Nagel 1974/1991:423).

2. What we discover when we examine conscious states is that they have a qualitative character, and their character is a fact about conscious states. To illus-

trate this point, Nagel contrasts human mental states with bat mental states. Because bat perceptual systems are radically different than our own, their states have radically different sorts of qualitative character. The character of echolocation is nothing like the character of any human state.

3. Moreover, only by having bat states can one understand the nature of the qualitative character they have. Only by being a bat can one conceive of what it is like to be a bat, and so the facts about bat experience are only accessible from the bat's point of view. "Whatever the status [representable in language or not] of facts about what it is like to be a human being, or a bat, or a Martian, these appear to be facts that embody a particular point of view" (Nagel 1974/1991:424) (4) Finally, subjectivity is having a particular point of view. Nagel writes, "every subjective phenomenon is essentially connected with a single point of view, and it seems inevitable that an objective, physical theory will abandon that point of view" (Nagel 1974/1991:423).

Subjectivity creates an obstacle to physicalism, according to Nagel, because facts about conscious states can only be described subjectively. Physicalism requires objective, physical descriptions of mental phenomena and so will leave out facts that are only describable subjectively. Most arguments against Nagel have taken aim at this objection to physicalism, but I am more interested in what Nagel says about subjectivity. A summary of the four claims central to the move from qualitative character to subjectivity can be stated thus:

1. What it's like to have qualitative character is the essence of consciousness.
2. There are facts about what it's like to have qualitative character.
3. These facts are special because they can only be described through a point of view.
4. Subjectivity is having a point of view.

I have no objection to the first and fourth of these claims (at least for the argument at hand). In Sections B and C I will consider the second and third claims, respectively, arguing that subjectivity provides a special route to facts, rather than a route to special facts as Nagel seems to hold.

4.2.B Nagel's funny facts

Taken by itself, the second claim above is not problematic. There are facts about what it's like to have qualitative character. Bat mental states are different from human mental states, and this difference is a matter of fact. But such facts, what Lycan appropriately calls *funny facts*,[117] assume a peculiar character in

the context of Nagel's argument. The first curious item is that the facts about what it's like to be a bat cannot be conceived of by humans; at least we have no idea how to form an adequate conception given our current resources. At best we can form a schematic concept of bat states by extrapolating from our own. Our schemas will necessarily be incomplete because we cannot supply the details of bat sensation. Whatever descriptions we concoct will be woefully inadequate to the qualitative richness of their conscious states.

> Thus we describe bat sonar as a form of three-dimensional forward perception; we believe that bats feel some versions of pain, fear, hunger, and lust, and that they have other, more familiar types of perception besides sonar. But we believe that these experiences also have in each case a specific subjective character, which it is beyond our ability to conceive. (Nagel 1974/1991:424)

According to Nagel, the inadequacy here is not simply the finitude of linguistic resources compared to the vast complexity and detail of conscious states. With enough time and care we could in principle coin sufficient terms for every nuance of qualitative difference.[118] Our language and concepts are limited not by detail but by access. We can form concepts about our own mental states because we have them, Nagel claims; we cannot form concepts about bat mental states because we do not have them. Facts about bat mental states are funny because they can only be conceived of from the bat's point of view.

Why should this be so when other things can be described from more than one point of view? You and I describe the same flower even if we generate our descriptions from opposite sides, and we can describe the same route even if you describe it from the point of view of a pedestrian (go left at the gas station) and I describe it by using maps (turn south on Highway 9). What is so special about bat mental states (or any conscious states) that description of them is limited to one point of view? Here Nagel invokes an argument already proven faulty: there is no appearance/reality distinction from the subjective perspective. For most things different particular points of view can be reconciled by appeal to an objective point of view because there is a difference between the way the thing appears and the way it really is. A rainbow can be described as it appears to me, standing in a field and looking at the horizon, but it can also be described scientifically, according to its constituent parts. Bat mental states, on the other hand, do not seem to have an objective nature; it seems wrong to talk about the way they really are apart from how they appear from the bat's point of view (Nagel 1974/1991:426). Nagel cleverly makes this point with rhetorical questions rather than straightforward argument. He asks:

> what would be left of what it was like to be a bat if one removed the viewpoint of the bat? But if experience does not have, in addition to its subjective character, an objective nature that can be apprehended from many different points of view, then how can it be supposed that a Martian investigating my brain might be observing physical processes which were my mental processes (as he might observe physical processes which were bolts of lightning), only from a different point of view? How, for that matter, could a human physiologist observe them from another point of view? (Nagel 1974/1991:425)

But the point must be made in order for the critical conclusion to follow:

> If the subjective character of experience is fully comprehensible only from one point of view, then any shift to greater objectivity – that is, less attachment to a specific viewpoint – does not take us nearer to the *real nature* of the phenomenon; it takes us farther away from it.
>
> (Nagel 1974/1991:425f, emphasis added)

If, however, bat mental states do have a real nature *apart* from the bat's point of view, then Nagel's conclusion does not follow. I see two ways in which the nature of bat mental states is separable from the bat's point of view.[119] First, a bat mental state *qua* representational vehicle is not fully comprehensible – not comprehensible at all – from the bat's point of view. Nagel clearly does not believe that the nature of mental states as representational vehicle is their *real* nature. As he asks, "how can it be supposed that a Martian investigating my brain might be observing physical processes which were my mental processes (as he might observe physical processes which were bolts of lightning), only from a different point of view?" (Nagel 1974/1991:425) Yet, if it is the case that physical states *are* mental states in the way that lightening bolts are, then surely the physical states have some claim to being the *real* nature of mental states. They have a metaphysical claim.

Second, and more to Nagel's point, a bat mental state *qua* bearer of representational content is separable from the object it is supposed to represent. Thus the intentional content of a representation may misrepresent its object. I see a yellowy-orange expanse but there is no yellowy-orange object or light source. What is the *real nature* of the phenomenon here? According to Nagel, it is the appearance of the yellowy-orange expanse, which can be known only from my point of view. But if this appearance is, as I have maintained, a way of representing yellowy-orange objects, then there is another sense to its real nature. That is, the real nature of the mental states is *as* a representation of yellowy-orange objects, and this nature is not known only from my point of view. In particular, I may not be in the best position to determine whether my

representation is doing its job, whether or not there is in fact a yellowy-orange object. So the objective perspective need not be any further removed from the real nature of mental states than is the subjective perspective. The real nature of mental states as representations may be assessed as well or better by the objective perspective. Alternatively, if mental states *don't* have a real nature separate from their appearance and their real nature is the appearance itself, then they cannot be merely intentional objects but suddenly become objects with their own metaphysical nature. Sense data. As we should have guessed, funny facts yield funny objects.

4.2.c Point of view provides a special route, not special facts

If bat mental states do not yield funny facts, then what makes them special is not that they are a unique sort of object – appearances. Instead we can say that bat mental states, or to get back to our target, conscious sensory states, are special in generating a unique sort of access or *route* to objects. Think of the way we represent objects as our *route* to them. The way we pick out objects is by representing them in some way or another. 'What it's like' for a subject to represent an object is one way of picking out that object.[120] The appearance to me of a red, shiny object on the table is one way I can pick out the apple. My appearances are special because they are unique to me; only I represent the apple in just the way I do. My route to the apple – via its appearance to me – is mine alone. In Part 3 I consider what exactly makes my (and your) point of view unique, but first we should finish with Nagel's argument. We can now see the flaw in Nagel's third claim. Rather than saying that a point of view reveals special facts, we should say that a point of view provides a special route to regular old facts. Subjectivity is special in giving the subject a unique way of representing things. But it does not follow that things represented in this way cannot be represented in other ways as well. I represent the apple as red and shiny from here, while to you over there it looks ripe and round. Though our points of view differ in various ways, we are both representing the apple.[121]

A final word about Nagel's views on objectivity may help shed some light on his subjectivity argument. To be perfectly objective, according to Nagel, requires "a transcendence of particularity and a transcendence of one's type" (Nagel 1979:209). The goal of objectivity is to represent things as they are in themselves, not as they are for some particular creature or type of creature. Since subjectivity is a particular mode of representation, a unique way of describing things, the subjective perspective will be left out of any perfectly objective description of the world. But this point has now become a logical truth,

posing no conflict with science. After all, science itself is a mode of representation and so is not perfectly objective. Nagel recognizes this point, noting that there is a continuum of objectivity which may have no endpoint where the completely objective nature of a thing is described (Nagel 1974/1991:425). The scientific mode of representation is more objective than the mode of representation from a single point of view because the scientific mode is more general. But this is only troubling if subjective modes of representation reveal meta-physical objects unavailable through scientific modes of representation. If not, then the fact that science fails to account for subjectivity is not a problem so much as a logical consequence of its more objective mode of representation.

4.3 Subjectivity as the view from here

Now that we have distinguished subjectivity from qualitative character, first person grammatical forms and appearances, what is left? Subjectivity is having a point of view, a unique perspective on the world that cannot be adopted by anyone else. As the previous section suggests, the uniqueness of a point of view can be cashed in terms of its special mode of representation. While several people may represent the same object and so share the same representational content, each person represents that object in a different way. If subjectivity is a mode of representation, then differences in individual modes of representation constitute the uniqueness in individual points of view. All of which sounds very nice until we ask what exactly a 'mode of representation' is and how we manage to get one. In this Part I will consider two theories that identify subjectivity with a mode of representation. The first is Lycan's account, developed in the context of his higher-order inner sense theory. The second is my own, based on the Gibsonian notion of an 'ecological point of view.' By examining these two theories we should get an idea of what a mode of representation might be and to what extent modes of representation might help solve the problem of subjectivity.

4.3.A Tokens in a language of thought

Lycan is rare among contemporary writers on consciousness for providing a positive materialist account of subjectivity, instead of either throwing up his hands in despair at the prospect or giving only a defensive response to Nagelian type objections. The account fits beautifully with his higher-order inner sense theory and accommodates several intuitions about subjectivity: its uniqueness,

inner aspect, and ineffability. In fact, I have nothing to say against Lycan's account except that, like his higher-order inner sense theory, it is a more sophisticated form of subjectivity than is necessary for sensory consciousness. My account describes a basic level of subjectivity developmentally prior to the level Lycan describes. Thus, I do not intend to supplant Lycan's theory so much as to supplement it.

Lycan begins by noting that some of the individuality in a point of view comes from the selectivity of perceptual organs. Because it is impossible to access or process all available perceptual stimuli, each view of an object or event will necessarily be somewhat different. Say that two people are looking at the same flower in a vase. By virtue of their different locations, each person will have access to a somewhat different set of stimuli than the other. One may see more of the flower's stem structure, the other more of the flower face. Likewise, each person will be attuned to different aspects of the flower. One likes the color, the other likes the pattern of the petals. So even if location differences were eliminated, each would notice some things about the flower and overlook others. Based on a lifetime of associations and acquired interests, such perceptual selectivity is as individual as a thumb print, or perhaps more so. A thumb print is permanent and therefore replicable, at least in science fiction. A brain print could theoretically be replicated, but the dynamic nature of synaptic interconnections means that differentiation would occur instantly. A person and her brain twin would be identical for only a fraction of a second before the flux of their separate systems would generate individual oddities.

Yet perceptual selectivity accounts for only part of the pull of subjectivity, according to Lycan. The uniqueness of subjectivity comes not only from what information a person gleans from her environment but also from how she uses that information. As Lycan observes, concepts have a dual nature. On one hand, concepts have truth conditions which determine their extension. On the other hand, concepts play a functional role in a person's cognitive system. In the proto-typical functionalist schema, concepts arise given certain kinds of input and cause certain kinds of behavior. What may be surprising is that truth conditions and functional role operate independently. Two concepts with different truth conditions may have the same functional role. These are the classic Twin Earth cases where 'water' plays the same functional role in the cognitive systems of Oscar and functionally identical Twin-Oscar, but 'water' refers to different substances on Earth and Twin Earth (Putnam 1975). Conversely, two concepts with the same truth conditions may have different functional roles. Here indexical cases make the best examples. I will react quite differently to the

information that PJD has won a scholarship than to the information that I have won a scholarship, unless I also know that I am PJD.

Lycan proposes that the particular dual nature of subjective concepts is the source of our most deeply rooted intuitions about subjectivity. Recall from Chapter 2 that the function of an inner sense, on Lycan's account, is to scan first-order sensations and produce second-order representations of them. Lycan takes these higher-order representations to be tokens in the subject's language of thought, so the output of an inner sense would be a mental word referring to the type of sensation scanned. What is important for the present discussion is that a mental word coined in this way would be utterly private because no one else could use this word to perform *both* its referential and cognitive role. It would be, as Lycan puts it, "a private name as well as semantically primitive, a name that only its actual user could use to name its actual referent" (Lycan 1996:60). As with Twin Earth cases, such a name could play the same cognitive role in two different people, yet refer to different mental states. For example, even if Lycan and I both use the word 'semantha' to refer to our intense cyan sensations, his use would necessarily refer to his sensation and my use would refer to my sensation. Because our inner scanners are wired only to our own first-order states – they are wired to no one else's states and no one is wired to ours – they produce mental words that refer uniquely to those states. Though 'semantha' plays the same role in our cognitive economies, the reference differs. As with indexicals, on the other hand, two names with the same reference could play different cognitive roles. 'C-fibers' and 'semantha' might both refer to my intense cyan sensation, and yet I could fail to realize 'c-fibers' are 'semantha.' Only my use of 'semantha' can play both its proper referential role (by referring to *my* sensations) and its proper cognitive role (*as* my sensations) (Lycan 1996:60).

Mental words are not logically private, however. Were another person wired to the same first-order states, she too could refer to them and would think of them as her own (Lycan 1996:172, Note 17). Yet even in this case, I would argue that a syntactical difference remains.[122] Imagine Oscar and Twin-Oscar hooked up to the same first-order states. Even if both used 'semantha' to refer to the numerically same cyan sensation, the word would need to carry the sense of 'my own' cyan sensation in order to play the appropriate role in each cognitive system. The owner to whom the sensation is ascribed would differ for Oscar and Twin-Oscar, presuming that they are two different people. The pronominal role of higher-order representations requires that mental states be ascribed to one's *self* in some way or another,[123] so no two persons could use the same higher-order representation to both refer to the same mental state

and to ascribe it to the proper self. As Lycan says, "the self-ascriptive combination of syntactic type and referent is invariably unique" (Lycan 1996: 60). Inner scanners produce representations of first-order states that are exclusive to the scanner's owner.

The exclusivity of second-order representations provides a potent explanatory tool in accounting for some of the most puzzling aspects of subjectivity. A subjective point of view is unique because each person has a different set of scanners and ascribes the sensory states scanned to herself. Only she can refer to her sensations by using the representations produced by her scanners. Other people can refer to her sensations through various means, but they cannot refer to her sensations in the way that she does. Also, a subjective point of view has a special inner aspect, a kind of inaccessibility to objective evaluation. Because the subject is the lone possessor of the scanners necessary to produce subjective representations and the representations thus produced play a unique cognitive role for the subject, they constitute a form of representation unavailable to anyone but the subject. It follows that such subjective representations are also ineffable. Though the referent of my second-order representations is describable in many ways, my special subjective way of describing the referent is unlike that of any other representational system. No public language terms represent my sensations in the same way they are represented in my own language of thought. Public terms are not semantically or syntactically similar to my own subjective representations, so public terms simply cannot represent my mental states the way I do (Lycan 1996: 64).

Given the power of Lycan's account, it may be tempting to adopt the higher-order theory of inner sense purely on the basis of its explanation of subjectivity. But such a move would be premature. We saw in Chapter 2 that the higher-order inner sense theory targeted a more sophisticated form of mental state than conscious sensory states, and the same can be said of the higher-order account of subjectivity. The pronominal aspect of higher-order representations does indeed account for the uniqueness and ineffability of our own ways of representing our mental states. But what of those creatures incapable of representing their own mental states? Humans arguably do not develop this talent until the age of 3 or 4, and there is no indication that animals introspect. Yet it seems reasonable to ascribe some form of subjectivity to them. Nagel's bat, for instance, seems to have a special way of representing the world, a unique point of view that requires explanation. Even if the bat is incapable of representing its own mental states, there is still a sense in which its first-order mental states themselves endow the bat with subjectivity.

What follows is a description of subjectivity that is meant to serve as a groundwork explanation for the sort of problems raised by Nagel. Because Lycan's account is framed in terms of conceptual differences between the way the subject represents her mental states (using private, pronomial names) and the way others represent them, his description of subjectivity directly addresses Nagel's worries about understanding the conceptual framework of the bat. If, however, the bat is incapable of introspection, then Lycan's explanation will not serve as a complete account. My proposal, therefore, is to provide a more basic description of subjectivity on the basis of *egocentric maps*. The idea stems from Nagel's suggestion that we develop an 'objective phenomenology' not based on empathy or the imagination. I agree with Nagel that "it should be possible to devise a method of expressing in objective terms much more than we can at present, and with much greater precision" (Nagel 1974/1991:427). In the next two sections of this Part I will try to express in objective terms a sense of subjectivity applicable to the bat and other non-introspective creatures by using Gibson's insights into the relation between subject and object.

4.3.B The reciprocality of subject and object

In his work on visual perception, psychologist J. J. Gibson realized that several faulty assumptions about vision had generated numerous theoretical puzzles. The most problematic assumption is what Gibson called "snapshot vision" (Gibson 1986:1). Researchers assumed that the eye functions as a sort of camera where each saccade generates a kind of picture of the scene. The notion of a retinal snapshot incorporates information about external qualities such as color and form within the internal processing system of the organism. On this view, the epistemic gap between subject and object is bridged by producing a literal copy of the external world inside the head of the organism. But as Gibson noted, the bridge is bought at enormous cost. As many or more questions are raised by snapshot vision as are answered. How is the information on the snapshot viewed? Is there a homunculus inside the brain that looks at the picture? We know that the retinal image is a reverse projection of the external environment. How is it re-inverted to look right side up? If each saccade produces an isolated snapshot, how does vision produce the impression of the environment as unified and 3-dimensional?

According to Gibson, one of the problems with the traditional view is its isolation of the eye as the single instrument of vision. On Gibson's ecological theory, the eye operates in coordination with the head and body to determine what sorts of resources, or *affordances*, are available in the environment (Gib-

son 1986:18). The goal of vision is not to get a picture of the world but to facilitate action in it. The subject is no longer removed from the object, requiring vision to bridge the gap, she is now centered within an environment, using vision as a tool for exploration. The organism depends on and defines an environment (Gibson 1986:8). Without an environment, no organism would survive, and conversely, the environs of an organism are what define an environment. Everyday perception presents an everyday world of surfaces which form such useful features as horizons, objects, openings and closures. No creature perceives an abstract geometric space of planes and lines, so perception theory defined exclusively in terms of these features will fail to explain how normal creatures perceive (Gibson 1986:54).

One of the consequences of placing the organism within an environment and *letting it move around* is that perception becomes more complicated. Eye, head and body movements must be included in the description of perceptual processes. Gibson ingeniously accommodates the locomotion of the organism by arguing that perception involves information about both invariant and perspective structure. Information that specifies stable environmental surfaces and information that specifies locomotion are reciprocal; each sort of structure specifies the other (Gibson 1986:75). For example, an organism can only determine whether it is moving past a car or a car is moving past it if it has information both about the car and about itself. If everything in the environment including the car is changing position relative to the organism, then the organism is in motion. If only the car position is changing relative to structures surrounding it and everything else around the organism is stable, then the car is moving. There are complicated cases where it may be difficult to assess what is moving relative to what, as when one is inside of a moving vehicle, but by and large information is available so that stability and movement can be specified relative to each other.[124]

Noticing the reciprocality of organism and environment led Gibson to the insight that subject and object are not separated by an epistemic gap but are two complementary aspects of every perceptual process. In the same ways one gains information about the world, one also gains information about one's location, body position and movements. Gibson writes:

> The supposedly separate realms of the subjective and the objective are actually only poles of attention. The dualism of observer and environment is unnecessary. The information for the perception of 'here' is of the same kind as the information for the perception of 'there,' and a continuous layout of surfaces extends from one to the other. (Gibson 1986:116)

According to Gibson, examination of the relations among surfaces is sufficient to determine the structure of both the organism and the environment. 'Here' identifies a unique set of surfaces that is defined by its relation to me, specifically its relation to my *nose*. The leftmost border of the right eye and the rightmost border of the left eye meet at the nose, making the nose "an absolute base line, the absolute zero of distance-from-here" (Gibson 1986: 117). 'There' also identifies a unique set of surfaces that is also defined by its relation to me. The only difference is that 'here' is occupied by me, 'there' is not.

Gibson speaks of subjective and objective perspectives as "two sides of the same coin" (Gibson 1986: 76). I believe this claim is the first step away from the mysterious, spirit world of mythical subjectivity toward a coherent, materialist account of the subjective perspective. Yet it is only a first step. Gibson's primary goal was to describe how an environment is structured so as to be knowable to an organism. But remarkably little is said about how the organism manages to exploit the information that the environment makes available. Gibson says that information pickup is improved when the perceptual system is "attuned" or "sensitized" (Gibson 1986: 254). Exactly what changes in the system constitute this attunement, however, remain obscure.

4.3.c Egocentric maps

Because my goal is to investigate the effects of the environment on the organism, specifically the subjective aspect of information pickup, it is necessary to consider how a perceptual system might be 'attuned' to environmental stimuli. As I shift attention to the subjective rather than the objective pole of perceptual information, I will make use of cognitive terminology, such as 'representation' and 'egocentric maps', eschewed by Gibson. As I will use them, these concepts seem to me to be fully compatible with ecological psychology, but it should be clear that Gibsonians may disagree about the extent to which the following amendments are welcome.

At the very least, I agree with Gibson that perception must determine the structures of both organism and environment simultaneously. The perceptual system identifies structures by extracting information from the sensory array that is sufficient to specify variant and invariant aspects of subject and object. Normally, however, we only notice the aspect of information that specifies objects. Because information about the environment is crucial to basic survival, attention is usually focused objectively rather than subjectively. What Gibson recognized was that our focus on the objective aspect of information obscured the fact that perception must simultaneously provide information about the

subject. As I move through the forest, the information about where my head is located is as vital to avoiding collision as the information about where the branch is located. Most often, though, this subjective aspect of information is used instrumentally while the focus remains on the environment.

In theorizing about the subject, we shift attention to the aspect of information that specifies the subject. What sorts of information might satisfy? The location and locomotion of the subject are certainly critical forms of information and more will be said about them shortly. More difficult to describe is information that determines whether a subject can use an object. As Gibson would put it, the subject perceives some kind of affordance yielded by the object. It does me no good to see the glass full of water on the table unless I also see that it is within my grasp.[125] What we need, then, is some description of the information that specifies notions such as 'within my grasp', or more generally, 'obtainable from here.' In other words, we need a description of *egocentric* information, information that specifies the subject.

Egocentric space, in contrast with *allocentric space*, is defined in relation to an organism within that space, according to the organism's ways of navigating and using that space. Allocentric space is defined in relation to no organism in particular. Objects in allocentric space are located with respect to one another, without reference to use by an organism. By looking at how egocentric space is determined, we may gain a better idea of what sort of perceptual information specifies the subject at its heart. John Campbell usefully describes egocentric space as a composite of coordinate systems referenced to the organism. These *egocentric frames of reference* are spatial positions drawn in relation to loci on the body (Campbell 1994:8). A person probably utilizes many egocentric frames of reference. In the way that Gibson specified vision in relation to nose, grasp could be specified in relation to the shoulder or torso, and audition could be framed from a base point midway between the ears.[126] By integrating information from these various egocentric frames, a person could then create a map of the environment specified in relation to itself.[127]

Campbell notes that egocentric coordination is essential to navigation, even when the coordinate system appears to be determined allocentrically. For example, rats in a water maze will learn to find a platform submerged in opaque liquid by using landmarks around the pool. Although the rat identifies the location of the platform in relation to the environmental landmarks rather than in relation to itself, the rat must triangulate the platform location with its own present location in order to swim to the platform. To know which way to swim, the rat has to determine where it is in relation to the landmark as well as where the platform is in relation to the landmark (Campbell 1994:21). Only at higher

levels of development do creatures acquire a thoroughly allocentric conception of space without even the sort of implicit egocentric reference necessary for navigation. But once we have such an allocentric conception, it becomes more difficult to recognize the egocentric aspect of perceptual information. A purely allocentric conception of space has enormous theoretical benefits. Allocentric space is the sort of space we need to do physics and geometry, to determine the relations between molecules, and to plot the stars. Of course egocentric space is necessary to everyday, non-conceptual activity; it is the sort of space we use get around. Yet we need not conceptualize the skills involved in utilizing egocentric space, nor need we even notice the ways egocentric space is used in navigating the environment. Egocentric space comprises the background information that makes perception and action possible, but it need never move into the foreground of attention or thought.

This is not to say that we *cannot* conceptualize egocentric space or that doing so will not bring theoretical benefits as well. Quite the contrary, I believe egocentric space is the place to find the information that specifies subjectivity. Think of the bat who cannot introspect and so cannot possess the subjectivity described by Lycan.[128] The bat's task while flying is to coordinate information from various frames of reference into a map of the environment. This map must be at least partly egocentric in order for the bat to determine its own location relative to locations of interest on the map. Note that it is insufficient for the bat to have a purely allocentric map of the bird's-eye view sort used on tourist information signs. Even with some marker labeling the location of the bat within the map – You are here – the bat must still be able to coordinate its own directional system with the directions given on the map. In effect it must place itself within the map to see that 'up' on the map means 'forward', 'right one inch' means 'right 100 feet'. The allocentric map must be interpreted egocentrically for it to guide action. An egocentric map, coordinated by front/back, right/left, up/down, calibrates locations in the environment relative to the bat. As the bat moves through the environment, the coordinates of environmental features on its egocentric map change accordingly.

Gibson argued that information about the structure of the environment could be extracted from information about changes in the coordinates of such features as edges (Gibson 1986: 72–75). An edge that remains uniformly distant despite changes in organism motion specifies a horizon, for example. The subject would also be specified by coordinate system changes but not in the same way. If an egocentric map is calibrated relative to the perceiver, her coordinates on the map never change. The perceiver is always at the locus, the zero point relative to which the map locations are calibrated. Therefore, the map

provides no direct information about the subject, precisely because the subject is not a positive feature of the map. Coordinate changes in environmental features specify the subject *indirectly*. When all the features on the coordinate system change at once, for example, it is likely that the subject is moving rather than the environment. When only a few coordinates change while most others remain stable, then the subject is probably at a fixed location.[129] Positive features of the subject – moving, at rest, 5 feet tall, upside-down – can be derived from the calibration procedures that form an egocentric map of the environment. But on the map itself, the subject is represented only implicitly as that which occupies the locus point. The locus is the point of view relative to which the map is coordinated. Subjectivity, on this way of thinking, is a matter of occupying the locus of one's egocentric map.

Suppose this account of subjectivity is true, how does it relate to sensory consciousness? The first thing to note is that a creature could be a subject on this account without having (or being capable of having) conscious sensory states. Unconscious sensory states are coded egocentrically and so could form an egocentric map with the sort of subjectivity I am describing. No conscious sensory states are necessary. Therefore, there can be no claim that the mystery of sensory consciousness has been surreptitiously shifted to the subject of consciousness. Since subjectivity is logically prior to sensory consciousness, on my view, nothing about the nature of the subject can explain consciousness.

Second, this account explains the subjectivity of sensory consciousness for non-introspective creatures like the bat. As coordinated representations composed of sensory states that per hypothesis code information egocentrically, conscious sensory states must also code information egocentrically. And so it seems phenomenologically as well. My conscious sensory states present the world as from here, cup to my right, window to the left, ticking clock behind. Conscious sensory states are subjective in the sense that they represent the world by using a coordinate system calibrated relative to the subject. Bat conscious states must be also calibrated according to a subject-relative system, and being the locus of this egocentric system accounts for the bat's subjectivity.

Third, conscious sensory states exhibit a coordinate system unique to that subject. Even if it were possible for two people to share the same set of perceptual systems, as in the hypothetical case of Oscar and Twin Oscar wired to the same sensory states, coordination of the information from these systems would nonetheless produce two numerically different maps. The coordinate systems would be identical – the two people would have equivalent egocentric maps – but the systems would be instantiated in different representational schemes. There would be two tokens for one map type and so two token locus points

and so two subjects. One might object that the coordinate systems must be based on the same locus point in order to be equivalent. But if the locus point is the same, then there could not be two token maps, since separate token maps require separate token loci. Thus there would be only one subject, not two.

Interestingly, the objection highlights an important ambiguity in the nature of a locus point. In one sense, the locus is the same for the two maps. Both are calibrated relative to the same position in egocentric space. Nonetheless the maps are numerically different and as such have numerically different loci. In terms of calibration systems, there is one locus; in terms of representation systems, there are two. So, two persons who shared the same perceptual system would have the same point of view in the sense that their egocentric maps would be calibrated relative to the same point. But they would differ in point of view in the sense that they would each be the possessor of a numerically different map with a numerically different locus. This result satisfies the paradoxical sense that subjectivity is essentially unique yet others can share 'a point of view' to a certain extent. Though two people cannot share the numerically same egocentric map, our maps can be calibrated more or less according to the same point. Filmmakers exploit this perceptual relation when they film from the perspective of a particular character. The technique allows audiences to assume the subjectivity of the character in a limited way by 'seeing through her eyes.'

A review of the results of this section reveals an intriguing if contentious description of subjectivity. A subject is the locus point of an egocentric map that coordinates perceptual information about features of the environment. Conscious sensory states are subjective because they likewise code information egocentrically. An egocentric map is numerically unique to the subject that produced it even if by weird science it is qualitatively identical to the egocentric map of another.

4.4 Deflating (and re-inflating) subjectivity

Is this account of subjectivity deflationary? In a sense it is. I have suggested that 'subjectivity', for all its vaunted mysteriousness, is no more than having a point of view, which in turn is no more than being the locus of an egocentric map. How terribly empty and disappointing this result is. Surely there must be more to subjectivity than I have admitted. The charge of concept deflation is sneaky because it is almost always true when some phenomenon is explained that formerly was regarded as mysterious. Explanations suck the mystery out

of a concept, and contrary to our vision of ourselves as knowledge-loving creatures, we are all enchanted by mystery. Our resistance to explanation is even more strong when it comes to mysteries about ourselves. Interior, inaccessible phenomena such as conscious sensory states and subjectivity comprise our private domain, a realm in which each of us is ruler with kingly privilege. To say that subjectivity is merely a locus point on an egocentric map is to strip bare our domain completely.

To soften the impact of this bleak portrayal, keep in mind that I have described only the most basic level of subjectivity. As we develop mentally, we develop resources for filling out our sense of ourselves as subjects. The most important of these resources is our ability to represent our own mental states. When we discover around age 3 or 4 that we sometimes represent the world truly and we sometimes represent the world falsely, we learn to view our mental representations *as* representations. This conceptual development opens up a whole new world for investigation, the inner world of mental representation. The non-introspective creature can only specify features of itself indirectly, by inference from information about environmental features. (I must be moving because everything in the surround is changing coordinates.) The introspective creature, on the other hand, can specify features of itself directly by representing its own mental states. This is, at any rate, the claim of Lycan's higher-order inner sense theory, and I find it persuasive.

The task of describing how the higher-order inner sense theory would change as a theory of introspection rather than one of sensory consciousness is a project best saved for another day.[130] But I will take a moment to describe how Lycan's account of subjectivity might connect with the account I have given. On Lycan's view, when I acquire the ability to represent my mental states, the resulting higher-order representations will be private and ineffable due to the pronominal aspect of subjective representation. My description of subjectivity as the locus of an egocentric map helps explain this pronominal feature.

Lycan argues that subjective representations are pronomial because they have a unique cognitive role and a unique referent. In Lycan's story my inner sense scans my sensations and produces higher-order representations of them. So a good bit of the ability to represent my sensations as mine comes from being wired to only one set of sensations, my own. But it is not good enough to simply demonstrably refer to *those* sensations and in fact always manage to pick out my own. I must represent them as *mine* in order for them to play the appropriate cognitive role. For subjective representations to count as *mine*, they must function inferentially and computationally as mine. What does this mean? It could mean the representation is explicit as when one says "Hey, that's

my foot you're stepping on." But more often the role is implicit; one simply acts on the basis of the ownership as, for example, by removing one's foot from further harm. My suggestion is, then, that egocentric maps form the basis of the ownership relation required for one to act in this way. At the most basic level, representations function inferentially and computationally as mine if they are coded egocentrically. *My* sensations are those coded relative to the same locus point, viz. the subject of those sensations. The locus point of my egocentric map serves to indirectly identify which sensations are mine: they are the sensations coded relative to me.

I, of course, am nothing more than a bare subject at this point, the locus relative to which sensations are coded. I am simply *here* relative to all the features that fill my environment. But as I begin to attach things like sensations to *here* when I acquire the ability for higher-order representation, I become less bare. Throbbing pain, fear, pleasure become located as *here*, not just as indicators of damage, danger, donut *there*. When I learn to represent my sensations *as* my sensations, rather than merely as representations of external features, I learn to represent myself as more than a bare subject. I become the bearer of various and interesting sorts of representations.

Unfortunately, I can only gesture wildly at this marvelous prospect. The goal of this chapter is much more modest: to show that the mystery of consciousness cannot be shifted to the subject of consciousness. At root, subjectivity is nothing more than the locus point of an egocentric map. As with theories of consciousness, when we take apart our pre-theoretical notion of subjectivity and carefully analyze its component parts, we find explanation within grasp. I have suggested that subjectivity be isolated from such associated notions as qualitative character, the grammatical use of first person terminology, and the difference between appearance and reality. All of these ideas are related to the concept of subjectivity, and their relations require deeper consideration. But, like a tightly knotted ball of string, the threads must be loosened from each other in order to identify each thread and trace its route through the intertwined mass. My suggestion is that when we tease apart the thread of subjectivity, we find a much thinner concept than the original ball would have led us to expect. It is this thin, thoroughly unmysterious concept that is at play in understanding the subject of sensory consciousness.

Chapter 5

Testing the theory

Now that we have the rough prototype for a new theory of sensory consciousness, it is time for a test run. The field of consciousness theory is currently so broad and varied that it is impossible to fully address all of the objections, thought experiments and alternate forms of explanation contained in the literature. In this chapter I will choose a few of the central problems relating to consciousness theory generally in order to place the second sense theory within some of these larger debates.[131]

5.1 Troubles with functionalism

The first two challenges are descendents from Ned Block's (1978/1991) worries about functionalism, of which the second sense theory is a variety. A functionalist proposes that mental states be explained in terms of their functional relations; what it is to be a mental state like pain is to have certain characteristic inputs, outputs and relations to other mental states. According to Block, functionalism is structurally flawed because there is no principled way to specify the right level of function to admit all and only creatures with mental states. If function is specified at the physical level – nerves and neurons and such – then creatures with a different physiology would be ruled out as mental creatures on an apparently arbitrary basis. Simply because you grew up on Mars or in an AI lab, and therefore have silicon synapses and a cyber cortex, you could not have mental states, no matter how well you behaved. This is pure chauvinism, Block says. Creatures physically unlike us might very well have mental states, so an adequate theory of mentality should accommodate them (Block 1978/1991:461). On the other hand, we musn't go too far. If function is specified at a broader level of psychological function, in terms of abilities such as language processing, then structures without any plausible claim to mentality will nonetheless count as having mental states. In his famous example Block proposes we interconnect all the people of China to instantiate the functional relations of a human brain. Supposing we could do this, Block argues,

these mere functional relations wouldn't endow the people of China with mental states except in the sense that each individual Chinese person has mental states. There wouldn't be an *additional* set of mental states formed by the functional relations. Here, Block admonishes, functionalism is too liberal (Block 1978/1991:451). To avoid Block's troubles, a functionalist must show that the theory has successfully navigated between the Scylla of chauvinism, where no one but humans have mental states, and the Charybdis of liberalism, where implausible structures such as a network composed of the people of China do have mental states.

5.1.A Chauvinism

Though Block's objections are aimed specifically at functionalist theories of mental states, the problems of chauvinism and liberalism apply to the second sense theory because it defines sensory consciousness in terms of its functional relations. I have maintained that conscious sensory states are coordinated representations of the world at the present moment produced by a second sense. Given these broad functional relations, the second sense theory is more likely to fall into the trap of liberalism than to founder on chauvinism. So let's pass the easier obstacle first. First, the requirement of a second sense itself may be considered too restrictive. It may seem chauvinist to claim that the only way conscious sensory states (i.e., coordinated representations of the present moment) could be produced is by a second sense. Why couldn't there be are creatures whose conscious sensory states simply appear *ex nihilo* with no precursor mechanism whatsoever? On the second sense account, such scenarios are impossible because a second sense is causally necessary to select which sensory states are coordinated into conscious states and which are not. The selection and coordination function performed by the second sense is required to produce states that represent the present moment. Without this control over the content of conscious sensory states, the essential temporal content of those states is lost.

This reply raises the deeper concern that the entire way of setting up the problem of sensory consciousness is chauvanist. I began in Chapter 1 by distinguishing unconscious, conscious and self-conscious states in terms of the way humans undergo these phenomena. Others have conscious sensory states only if they have systems relevantly similar to ours, where the criteria for being 'relevantly similar' states are: (1) coordinated representations of the present moment, (2) which are produced by a second sense. These are taken to be necessary and sufficient criteria for sensory consciousness because they are neces-

sary and sufficient for our conscious sensory states. Isn't this anthropocentric definition a form of chauvinism? It may be, but I don't think it is either avoidable or objectionable. The theory proposes in the first instance to explain a phenomenon identified by and puzzling to humans in particular, the phenomenon of sensory consciousness. We begin with our own conscious sensory states and argue outward from there. Really, we have very little choice. Though I argued in Chapter 4 that the subjectivity of sensory consciousness does not pose an ontological problem for materialist, it does pose a methodological one. The best form of evidence we have for conscious sensory states is first-person reports. In tandem with neurological data from various sources and behavioral evidence, such reports give us access to an otherwise private phenomenon. As far as we know, we are the only creatures that can issue first-person reports, so we must begin our investigations with the phenomenon of sensory consciousness as it is manifest in our own mental system.

That said, it is worth noting that the theory does not *in principle* limit conscious sensory states to humans, and so is not chauvanist in the sense that troubled Block. Any creature – animal, infant, Martian or otherwise – has conscious sensory states if that creature has coordinated sensory representations of the world at the present moment produced by a second sense. Infants and primates certainly meet these criteria, and many mammals seem to as well. Lower down the phylogenetic scale things get more fuzzy, of course. Note, though, that creatures with a different physiology altogether are not ruled out of contention for the status of conscious creature on my account. Even if you grew up on Mars or in an artificial intelligence lab, and therefore have silicon synapses and a cyber cortex, you could still have conscious sensory states. Creatures quite different from us physically might very well have conscious sensory states, provided they have some sort of second sense mechanism that selects and coordinates sensory representations into a representation of the present moment. Because the theory is not committed to a particular physiological instantiation of its functional organization, a wide variety of systems might manifest sensory consciousness on this account.

5.1.B Liberalism

On the other hand, we musn't go too far. The functional definition of sensory consciousness must sail wide of chauvinism, but it cannot go too wide lest it fall prey to liberalism. A theory that is too liberal becomes swamped by structures that are not plausible candidates for sensory consciousness. In the case of the second sense theory, it may seem that a structure could possess represen-

tations of the world at the present moment produced by a second sense and yet there would not be anything it is like to be that structure. For example, a weather computer might coordinate input from a satellite feed, barometer and thermometer to form a coordinated representation of the weather at the present moment. Yet no one would say that the weather computer was therefore endowed with sensory consciousness. If there is one thing we ought to be able to say with certainty in this contested arena of consciousness theory, it is that there is nothing it is like to be a weather computer.[132]

Indeed, one could argue that no explanation of sensory consciousness is possible, because there always remains the appalling possibility that we might build (or find) a creature that satisfies the conditions of explanation yet lacks conscious sensory states. This sort of philosophical zombie differs from the horror film variety in that philosophical zombies are not wide-eyed, stumbling and disfigured. What makes them appalling is that they seem just like us in every respect, except there is nothing it is like to be them.

The first step toward diffusing zombie worries is to note how difficult it will be to fulfill the necessary conditions for sensory consciousness. For one thing, the representations coordinated into conscious sensory states must be genuine mental representations. While it is a matter of debate what constitutes a genuine mental representation, it is generally agreed that conventional forms of computer processing are insufficient.[133] The average computer is capable of representing only to the extent that it has been programmed to represent. It has no ability to generate its own representations; it is limited to the language of its program.[134] So the weather computer used by the local meteorologist can be ruled out as a counter-example. But suppose we eventually do build a device capable of genuine mental representation, whatever this might entail. On my view its representations would still be unconscious unless the further conditions for sensory consciousness are met. So next we build a device with genuine mental representation and add a coordinating mechanism. What reason is there to believe that this new device now has conscious sensory states?

The answer lies, as one might expect, in the nature of the coordinating mechanism. To be capable of producing conscious sensory states, the coordinating mechanism cannot simply throw together whatever representations are available. In order to count as a second sense, the mechanism must be able to select and coordinate representations in a particular way. Conscious sensory states are, according to the proposed theory, coordinated representations of the world at the *present moment*. This means that the selection process of the coordinating mechanism must include a temporal component as part of the resulting representational content. It is insufficient for the coordinating mech-

anism simply to coordinate the representations now occurring. Our unconscious sensory states can be so coordinated and yet remain unconscious. To be a second sense, the mechanism must coordinate representations so as to produce a representation *of* the present moment. It is far from clear to me how one might build such a mechanism. As noted, we would first have to determine how to build a device capable of genuine representation. We would then need to consider exactly what is included in 'the present moment.' In Chapter 3 I argued that the sensory states included in a representation of the present moment are the best approximation of the world, coordinated toward accomplishing a particular sort of task. Thus, not all sensory representations occurring at a given moment will be incorporated into a representation *of* that moment. Selection, and even confabulation, is critical to the function of a second sense. To build a second sense, therefore, would require careful attention to the processes of selection and coordination accomplished by this unusual and sophisticated mechanism.

If such a mechanism could be built, then on my view we could conclude that there is 'something it is like' for the system we have built. It would have sensory consciousness. World-representation, executive planning, perspectival navigation – all describe important and perhaps necessary features of sensory consciousness. Nonetheless, these features are not sufficient. The essential element to conscious representation so often missing in physicalist theories of consciousness is the ability to represent 'now'. So coordinated representations, even if they are genuine sensory representations, are not sufficient to count as conscious sensory states unless they include this critical temporal aspect.

Finally, I have argued that conscious sensory states submit of degrees, coordinated relative to the task at hand. If one is engaged in a highly specific task, well-coordinated representations are needed. Whereas loosely coordinated representations are sufficient for more general tasks. An adequate reproduction of a second sense mechanism must be able to adjust its coordination strength along a similar continuum. Furthermore, if a representational system has very limited coordination ability and sensory systems quite unlike ours, its conscious sensory states are likely to be phenomenologically different than ours. As with Nagel's bat, we may find it difficult to imagine what it's like to be such a thing. We should not suppose, however, that our inability to imagine the conscious sensory states of such a strange thing means that it has no sensory consciousness. Such a supposition, given the above requirements are satisfied, would be an extreme form of chauvinism.

And so the theory skirts the parallel dangers of liberalism and chauvinism. It will not be simple to develop a system capable of genuine representations,

much less the particular form of representation required for sensory consciousness, coordinated representations of the present moment. These stringent requirements severely limit the number of structures that count as conscious sensory states, effectively ruling out implausible structures such as the weather computer. Yet the requirements are sufficiently loose that physiological structures markedly different from ours could satisfy them. Animals, infants and Martians all might have conscious sensory states, provided they have coordinated sensory representations of the world at the present moment. The ship sails smoothly past the Charybdis of liberalism and the Scylla of chauvinism.

5.2 The hard problem

More dangers await, however. In the previous section the theory successfully avoided liberalism by setting rigorous conditions for sensory consciousness. Yet another sort of zombie remains a threat: the logically possible zombie. It seems conceivable that a creature could have coordinated sensory representations of the world at the present moment produced by a second sense and yet there could still be nothing it is like to be that creature, contrary to the second sense theory. The central worry here has been characterized by David Chalmers as the Hard Problem of consciousness. We cannot define sensory consciousness in functional terms – such as coordinated representations of the present moment produced by a second sense – because it would always be possible to conceive of a structure that instantiated those functions without there being something it is like to be that structure.

Many other sorts of mental function form the relatively Easy Problems of consciousness, since we know how to go about explaining them. Some of the easy functions are:

> the ability to discriminate, categorize, and react to environmental stimuli; the integration of information by a cognitive system; the reportability of mental states; the ability of a system to access its own internal states; the focus of attention; the deliberate control of behavior; and the difference between wakefulness and sleep." (Chalmers 1997a: 10)

By contrast, the Hard Problem of consciousness is explaining the 'what it's like' of conscious sensory states, their felt quality (Chalmers 1997a: 10). The Hard Problem is hard because this experiential aspect of consciousness cannot be defined in terms of function. After all of the functions of consciousness are explained, there is still the further question of why there is something it is

like to perform these functions, why they are accompanied by "experience" (Chalmers 1997a: 13). Chalmers considers two attempts at functional analyses of consciousness (one by Crick & Koch [1990b] and the other by Baars [1988]) and concludes:

> At the end of the day, the same criticism applies to *any* purely physical account of consciousness. For any physical process we specify there will be an unanswered question: Why should this process give rise to experience? Given any such process, it is conceptually coherent that it could be instantiated in the absence of experience. (Chalmers 1997a: 18)

One issue in the Chalmers challenge is the relation between the definition and explanation. Must something be *defined* in physical terms in order to *explain* it in physical terms? Though water is not defined as H_2O, this is no threat to the explanation of water in physical terms. So why should the explanation of consciousness be treated differently? According to Chalmers, even a posteriori identities like H_2O = water are logically derivable from physical theory. The low-level physical facts imply the higher-level facts as a matter of conceptual necessity. We begin with water and what it does: clear, colorless liquid, boils when sufficiently heated, freezes at 0°C, etc. Then, through the course of scientific investigation we develop a chemical theory of molecules and atoms and their interactions. Given this chemical theory we can then *deduce* that H_2O molecules will exhibit the features commonly identified as 'water.' Hence, H_2O = water. The same is true for genes and DNA and all other a posteriori identities. But consciousness cannot be derived from any physical theory, so to posit an identity between consciousness and, say, the brain relies on an "explanatorily primitive link" (Chalmers 1997b: 390). Other scientific identities have been earned by showing how the two sides of the identity are identical, how the physical substance or function has all the properties required to be water or genes or whatever. Without a similar demonstration, Chalmers maintains, the suggestion of an identity does not explain consciousness by virtue of an identity, it posits an identity in place of an explanation (Chalmers 1997b: 390).

Chalmers has a point here. At this point in consciousness theory any suggestion of an identity between conscious states and brain states is more an exercise in speculation than scientific proof.[135] If one of these theories is true, then of course the identity is a posteriori necessary. But no one has provided the kind of explanation required to support an identity between conscious states and brain states in the way that chemical theory supports water = H_20 or biological theory supports genes = DNA. I do not think, however, that Chalmers

can conclude that such an identity must therefore be a fundamental physical law (Chalmers 1997b: 390f). While fairly brute correlations are all that support the identity at the moment, further research will no doubt show ever more intimate connections between the phenomenon of sensory consciousness and possible physical instantiations. As these connections are forged and physicalist theories of sensory consciousness develop, we can expect that our concept of sensory consciousness will be reconfigured. Paradigm shifts characteristically involve such definitional changes to accommodate new insights and theoretical interconnections.

Furthermore, the claim that all a posteriori identities are logically derivable from physical theory is contentious. As Block and Stalnaker (1999) argue, the sort of a priori derivation Chalmers suggests fails to go through even for the simple case of explaining the fact that water boils. Difficulties multiply when explaining biological functions such as reproduction. Chalmers and Jackson (2001) offer a complex reply that illuminates the epistemic flexibility of a priori entailment they envision. Nonetheless, the reply fails to address the critical question of whether all explanatory laws are reducible to physical laws. Certainly there is "*a* macroscopic description of the world in the language of physics" (Chalmers & Jackson 2001: 330), but it is not so certain that this description adequately captures all the laws of biology and psychology, let alone the laws of economics and sociology. In arguing that all identities are in-principle deducible, Chalmers and Jackson note that one could deduce that Mark Twain is Samuel Clemens from all of the physical, mental and social truths available.

> The subject will be in a position to know that there was an individual who was known to his parents as 'Samuel Clemens,' who wrote books such as *Huckleberry Finn* and the like under the name of 'Mark Twain,' whose deeds were causally responsible for the current discussion involving 'Mark Twain' and involving 'Samuel Clemens,' and so on. From all this information, the subject will be able to easily deduce that Mark Twain was Samuel Clemens, and the deduction will be a priori in the sense that it will not rely on any empirical information outside the information specified in the base.
>
> (Chalmers & Jackson 2001: 355)

While I agree with the conclusion, the inclusion of concepts like 'parents' and 'book' in the premises highlight the limits of a macrophysical description of the world in the language of physics. Neither of these terms is specifiable in the language of physics. Even though there is a description of me in purely physical terms, this description would not include the fact that I am a parent. Nor would a purely physical description of this book include the fact that it is a book.

What makes this a book and me a parent involve laws that are not implied by physics. Because Chalmers and Jackson often talk of 'natural phenomena' and take 'water is H^2O' as their primary example of a macroscopic truth, they may not mean to extend the claim of a priori entailment to social phenomena. If so, then we may be faced with a new Hard Problem once we have addressed the problem of consciousness.

A second issue raised by the Chalmers challenge is whether consciousness can be defined functionally. Is it always possible to ask the further question of why there is something it is like to have certain functional states? In one sense, it is possible. As I noted at the very beginning of this book, some terms are suffused with a cultural and moral mysteriousness that no definition or explanation could remove. For example, it may be that 'life' can be defined in terms of adaptive behavior and reproduction, and perhaps the vitalists would agree (Chalmers 1997a: 18). Still, it is conceptually possible, to me anyway, that cells and DNA and the rest could go about their business without resulting in *life*. I know my share about the birds and the bees, but there still remains the question of how mere cell division results in a *live baby*. Maybe I simply don't understand the mechanisms well enough. But if ignorance is the only problem, then ignorance of the mechanisms resulting in sensory consciousness is likely the obstacle to an adequate explanation in this case too. When we just get the right theory of consciousness, then we will have just as satisfying (or unsatsifying) a sense of explanation as in the case of life.

So there is another sense in which it is not possible to ask the further question of why there is something it is like to have a functional state. I have offered a definition of sensory consciousness in functional terms. It is not open to the objector simply to say that I have left out the 'what it's like' of sensory consciousness. My claim is that 'what it's like' to have a conscious sensory state *just is* to have a coordinated sensory representation of the present moment, produced by the second sense. The theory claims to account for all the phenomena that seem to me to be associated with sensory consciousness. To challenge this account requires specifying what exactly has been left out. Just as Mendel's genetic theory was later supplemented with more specific theories about chromosomes, the second sense theory is open to a more detailed account of sensory consciousness. For example, further investigation into the possible mechanisms for producing conscious sensory states may reveal evidence in favor of the sort of hybrid higher-order/flat theory considered in Chapter 3. But just as it is not possible to ask why things that genetically reproduce and adapt are alive, it is not possible to ask why there is something it is like to have coordinated representations of the present moment. The mysterian about conscious-

ness is in the same position as the vitalist about life. A residual air of wonder – why life? why experience? – does not in itself count against the definition unless a reason can be given to believe the definition fails in some particular way.

5.3 On Rosenthal

The remaining two parts deal with challenges to fundamental features of the second sense theory. In this part I take up Rosenthal's claim that thoughts are necessary for conscious sensory states. Curiously, Rosenthal often presents his higher-order thought theory as simply the 'only alternative' to the higher-order perception view (Rosenthal 1991a:31, 1993a:158, 1993b:361). This modesty masks a deep theoretical commitment to concepts as essential for conscious sensory states. As I noted in Chapter 2, the notion of 'thought' central to the higher-order thought theory is less than clear, but concepts are definitely required. In support of this requirement, Rosenthal gives two intuitively persuasive reasons to think concepts are necessary to conscious sensory states. First, the acquisition of concepts adds complexity and subtlety to the content of conscious sensory states. Wine-tasting is the common example, "where conceptual sophistication seems actually to generate experiences with more finely differentiated sensory qualities" (Rosenthal 1991a:34). The higher-order thought explanation for this phenomenon is that the subtle qualities of the sensory state remain unconscious unless the higher-order thought is sufficiently subtle to represent them. "The degree to which we are conscious of differences among sensory qualities," Rosenthal explains, "depends on how fine grained the concepts are that figure in our higher-order thoughts" (Rosenthal 1991a:34). On this view, concepts not only seem to generate the conscious sensation of finely differentiated qualities, concepts actually do generate finer conscious sensory states. More of the sensory qualities are in fact conscious when one acquires more sophisticated concepts. Therefore, Rosenthal concludes, "it is unlikely that we can explain these observations except by a theory that appeals to higher-order thoughts" (Rosenthal 1991a:34).

Rosenthal's conclusion seems reasonable enough, but it rests on the sly shift from the way new qualities *seem* to become conscious with new concepts to the claim that new qualities *in fact* become conscious with new concepts. If the latter were true, I agree that Rosenthal would have a strong case for some form of cognitivist view. But why should we think that new qualities of sensory states are conscious due to new concepts? Why not hold, as the second sense theory claims, that the qualities were conscious all along, and the new concepts

merely helped in differentiating them? Acquiring new conceptual resources aids in many sorts of discrimination tasks, so it is reasonable to expect the same sorts of benefits regarding conscious sensory discrimination. Moreover, it may seem to us that new qualities become conscious with the acquisition of new concepts, but a contrasting intuition counters this one. That is, it also seems to us that our conscious sensory states are filled with more detail than we can conceptually comprehend. One contemporary film editing technique exploits these conceptual limits by bombarding the viewer with a rapid succession of images and sounds. The overall effect of this technique is to leave the viewer with a sense of rich audio-visual detail but only a general impression of what events have transpired. So at minimum these competing intuitions about our conscious sensory states cancel one another out, and we cannot depend on conscious sensory discrimination to be decisive in favor of a cognitivist theory.

The second reason Rosenthal gives in favor of the higher-order thought theory rests on what a person is able to report about her conscious states. Rosenthal begins with an analysis of speech acts. When one says something meaningfully and sincerely, one expresses a thought. Since a report that one is in some mental state is a form of speech act, it expresses a thought of some kind. Speech acts and the thoughts they express have the same propositional content, so a report that one is in mental state *S* expresses the thought that one is in mental state *S*. One could directly express the mental state in some form other than a report, such as through one's actions or tone of voice. An explicit report, however, expresses a thought about a mental state and does not directly express the mental state itself. Belief reports show this distinction most clearly. I can express my belief that *p* simply by saying '*p*' or I can report this belief by saying 'I believe that *p*'. My report expresses a different mental state than the belief, it expresses the thought that I believe that *p* (Rosenthal 1993c: 200–204). Thus far I agree. Though my brief review glosses over the intricacy of Rosenthal's argument, the central connection between a report about a mental state and a thought that one is in that mental state should be clear. And, as an analysis of the distinction between expressing and reporting mental states, I believe Rosenthal is right on target.

But then Rosenthal goes on to draw conclusions about the nature of consciousness on the basis of this distinction between expressing and reporting. He says,

> given that a creature has suitable communicative ability, it will be able to report being in a particular mental state just in case that state is, intuitively, a conscious mental state. If the state is not a conscious state, it will be unavail-

> able to one as the topic of a sincere report about the current contents of one's mind. And if the mental state is conscious, one will be aware of it and hence able to report that one is in it. (Rosenthal 1993c: 204)

Apart from the appeal to our intuitions about conscious mental states, Rosenthal seems to be making a technical point about reportability: a mental state is reportable iff it is conscious.[136] This claim draws a very strong connection between reportability and consciousness that I do not find intuitive at all. Though Rosenthal is careful to specify that the requisite language skills are required for reports, deficient language ability is only one problem that might render the contents of a conscious state unreportable. A creature might not have the necessary conceptual abilities to form thoughts about conscious states, so an inability to report these states would not imply that they are not conscious, unless thoughts are necessary to consciousness. But the connection between thoughts and conscious states is what Rosenthal wants to prove, so it cannot be used as a premise here.

Additionally, factors may intervene between the conscious state and a report about it. I have argued that there is a principled difference between Dennett's Stalinesque and Orwellian forms of confabulation.[137] There is a matter of fact about whether a person edits the representation of an event, such as the cutaneous rabbit arm taps, before it becomes conscious or afterward. If the Orwellian story is true, then people have the conscious sensory representation of three separate sets of taps and yet seconds later (falsely) report having consciously felt a continuous sequence of taps. If indeed this is possible, then a state's being conscious does not entail its being reportable. Moreover, the problem of confabulated reports counts against the implication from reportability to consciousness as well. A person may quite sincerely, but mistakenly, report having a sensory state that she does not have. There is no reason to believe that we have infallible access to our sensory states, and so no reason to believe we always accurately report our sensory states, conscious or otherwise. Unless one already holds that what constitutes a conscious sensory state is having a thought about it, the intimate tie between reportability and consciousness comes undone. While higher-order thoughts, that is, thoughts about sensory states, do seem to be closely linked with reports about sensory states in just the way Rosenthal argues, no such link can be used to support the higher-order thought explanation of conscious sensory states.

Therefore, there seems no reason to believe that concepts are necessary in order to have conscious sensory states. Not only is it difficult to define the conceptual abilities required, the reasons in favor of such a requirement are not

decisive. We can explain why it *seems* that new sensory qualities are conscious when we acquire new concepts without claiming that they in fact *do* become conscious. They were conscious all along, and the concepts merely aided in discriminating them. And we can accept the connection between reporting a sensory state and expressing a thought about that sensory state without accepting any connections at all between reports and conscious sensory states.

5.4 Dealing with Dennett

This final part is duly devoted to the philosopher who made consciousness a respectable topic for contemporary philosophical investigation, Daniel Dennett. With inimitable style and a dash of fascinating empirical research, Dennett has in many respects framed the terms of a solution to the problem of consciousness. I have already taken up several of his examples and test cases and will consider two more in the next two sections. The first is an objection Dennett raises to Dretske's distinction between 'thing-awareness' and 'fact-awareness,' a distinction I hold as well. Dennett argues that this is a distinction without a difference because there is no way to explicate 'thing-awareness' that does not involve an awareness of facts, or 'micro-cognitions' as Dennett calls them (Dennett 1994:513). The second section deals with an objection Dennett raises to higher-order theory but applies to my flat theory as well. According to Dennett, any theory that does not take first-person reports as criterial to consciousness is open to the "bizarre category of the objectively subjective – the way things actually, objectively seem to you even if they don't seem to seem that way to you" (Dennett 1991:132). Both of these objections stem from Dennett's conviction that consciousness is nothing more than first-person judgments and the reports they generate. In claiming that a creature can have conscious sensory states without a report, judgement or even a concept, my theory is directly opposed to Dennett's description of consciousness.

5.4.A The Camera Obscura argument

The first of the two objections is Dennett's most serious attack on a non-conceptual theory of consciousness such as the second sense theory. In distinguishing between 'thing-awareness' and 'fact-awareness' Dretske claims that is possible to have a sensory representation of something without any concept of the thing represented.[138] Dretske explains,

> A mouse in the kitchen . . . can smell, and thus have sensory awareness of burning toast but it will not (like the cook) be aware (i.e., believe) that toast is burning. The mouse will have little or no conceptual awareness of this event. It will smell the toast burning but not smell that it is, not smell it *as*, burning.
> (Dretske 1995: 10)

In contrast, 'fact-awareness' does involve concepts. Unlike the mouse, the cook does believe that the toast is burning because, unlike the mouse, the cook has and can apply the concept 'toast' and the concept 'burning'. 'Thing-awareness' can lead to 'fact-awareness', given the right epistemic conditions, but they are not the same. Similarly, I have argued that there can be concept-free sensory representations and appropriate composites of these could be concept-free conscious sensory states.[139]

Dennett denies Dretske's distinction, arguing that there is no such thing as a sensory representation without some form of cognitive uptake. The category of 'thing-awareness,' Dennett claims, is "an artifact of taking ordinary language too seriously. There is no important difference – no difference that makes a difference – between things nonepistemically seen (e.g., the thimble in front of Betsy's eyes before she twigs) and things not seen at all (e.g., the child smirking behind Betsy's back)" (Dennett 1994: 511). To me the difference between these two cases is as obvious as it is for Dretske. Just because an event happens to have no epistemic consequences does not mean it does not occur, nor does it mean the event is non-mental. On a teleo-functional account such as both Dretske and I hold, mental representations are defined in terms of their historical effectiveness, not by the consequences of a single event of representing. Less obvious is how to determine when a particular mental event has occurred in the absence of epistemic consequences such as a first-person report. As Dennett puts it, the problem is to distinguish between what is nonepistemically seen and what is visible. Nonepistemic seeing must be more than simply having visual information, because numerous devices process visual information without the ability to see. As an example, Dennett describes a precursor to film cameras, the camera obscura, which is simply a dark box with a pinhole opening. Light refracting through the pinhole is projected on the back of the box, displaying a visual copy of the scene in front. But it would be absurd to claim the box could 'see' the scene. In Dennett's view, the addition of film is "a step in the right direction", because the scene now leaves a trace. Still, a camcorder does not see (Dennett 1994: 512). For Dennett, really seeing requires some form of cognitive uptake. Unless seeing issues in what Dennett

calls 'microcognitions,' "we are stuck unable to tell the camera obscura from the genuine seer" (Dennett 1994:513).

In light of Dennett's Darwinian desires for continuity from mindless, mechanical processes to full-fledged narrative self-representations, his objection the idea of non-cognitive sensory representation is odd. Dennett's explanatory strategy, it seems, is to describe the mind in terms of cognition and then establish the Darwinian continuity from mind to mindless as successively degraded or precursor forms of cognition. So, in terms of vision, Dennett claims:

> What a genuine seer must do is somehow take in and "categorize" or "recognize" or "discriminate" or "identify" ... (each term stretched out of its ordinary field) ... or in some other way "judge" the presence of something (as a thimble or as something else). With such uptake there is seeing. Otherwise not.
> (Dennett 1994:513, ellipses original)

Dretske's category of 'non-epistemic seeing' fails to result in any sort of cognitive uptake, even of a rudimentary kind. Therefore, according to Dennett, it is no different than other non-mental processes such as exhibited by the *camera obscura* and the camcorder.

In the same vein, Dennett argues that consciousness depends on having a conceptual structure. Conscious states are ones that achieve what he calls 'cerebral celebrity' – they last long enough or have sufficient effects to be reportable (Dennett 1994:547). As Dennett succinctly states in his basic methodological principle, first-person operationalism: "If the subject can't report it, it isn't part of the subject's consciousness" (Dennett 1994:545). The necessity of reportability dovetails with Dennett's claim that concepts are required for consciousness.[140] Because concept acquisition is a matter of acquiring a new competence, it is reasonable to propose that, since consciousness is useful, it too involves acquiring concepts (Dennett 1994:550). Conscious states are useful because they perform a cognitive function; otherwise what reason would there be to bother with them?

While cognitive functions are the most useful mental functions, and perhaps the ultimate goal of all mental function is to aid cognition, there are reasons to suppose some mental functions are not explicitly cognitive. First is the problem of concept acquisition. Presuming some of our concepts are formed through sensory interaction – by means of demonstrative reference, for example – there must be information available in sensation to facilitate the formation of concepts. The sort of information available in the camera obscura might be sufficient, but in any case it seems that sensory processing goes beyond the one-to-one pictorial image produced by the camera obscura. In my

view, sensory processes go so far as to be representational; that is, senses have the function of representing properties of objects such as their color, sound, and texture. When things go wrong, senses can misrepresent. The blue tie *looks* green, is represented visually as green, but it is in fact blue. Sensory misrepresentation is different from miscategorization. I don't mistakenly *think* the blue tie is green, I may know perfectly well that it is not in fact green. I *see* the tie is green, but my visual sensation is misrepresenting the color of the tie.[141]

If it is true that sensory processes have the function of representing properties of objects, they could perform this function even if no concept was formed as a result. At some point concepts must be formed in order for sensory processes have long-term value. As Paul Churchland says, the mind is an 'epistemic engine' with knowledge as its fuel (Churchland 1979). This does not mean, though, that each individual sensory representation must have the function of forming a concept. Just as the function of cognition per se is to aid the survival of the creature (or the species, as the case may be), individual concepts may fail to perform this function. The concept might be faulty or, as so often happens, irrelevant. Likewise sensory representations could fail to result in concepts even if concept production is an important function of sensory representation. Moreover, there may be quite simple creatures capable of sensory representation but incapable of concept formation. Such creatures could use these representations to guide activity in the absence of the higher level abilities of individuation and reidentification.

I do not mean to suggest that the idea of non-conceptual sensory representation is uncontroversial. Dennett may limit the category of mental representation to conceptual representation alone and deny that there is any such thing as sensory representation. However, he will need more argument for this restriction than the supposition that cognition is the only valuable mental function. There is reason to believe our senses represent and could perform this representational function whether or not a concept (or even 'microcognition') was the result. And, if we take conscious sensory states to be coordinated sensory representations of the world at the present moment, they too might perform their function without resulting in the production of either concepts or reports.

5.4.B The bizarre category of the objectively subjective

One of the reasons Dennett is unlikely to accept the conclusion of my teleofunctional argument in the previous section is his worry about the "bizarre category of the objectively subjective" (Dennett 1991: 132). According to Dennett, the 'subjective' is constituted by first-person judgements,[142] so it is 'bizarre'

to apply the third-person standards of objectivity to the category of the subjective. How things seem to you is how you judge that they seem. This view is an immediate consequence of Dennett's first-person operationalism, which "brusquely denies the possibility in principle of consciousness of a stimulus in the absence of the subject's belief in that consciousness" (Dennett 1991:132). Sensory consciousness perfectly coincides with subjectivity in Dennett's view. To inject objectivity at this point would be to claim that a person might be wrong about how things seem to her, that some other person could correct her first-person judgements. And, while Dennett is the first to admit we are not always right about the way things are, he does seem to think we are always right about how things seem. The argument, reviewed in Chapter 3, against a distinction between Stalinesque and Orwellian revisions of perceived events essentially comes down to this point. There is just nothing more to seeming – no facts of the matter about the contents of our conscious sensory states – than what we judge. Since our reports are reports of our judgements, they too must be accurate about how things seem. Consequently, it makes no sense on this way of thinking to talk about how things really seem if they do not seem to seem that way. The contents of first-person reports comprise all there is to how things seem, there is no 'really seeming' to contradict them.

On the second sense theory, however, there can be such contradictions. If we take 'how things seem' to be the contents of conscious sensory states, then things can seem to be one way, and yet be reported as seeming to be a different way. They can be represented one way in a conscious sensory state and represented differently in a first-person report. Take the first sort of seeming to be 'how things really seem' and the second to be 'how things seem to seem' and when they conflict, we have the category of the objectively subjective. I do not find this category at all bizarre. Rather it is simply one more instance of the general post-Cartesian discovery that we are not nearly as knowledgeable about our own minds as we may have thought. Far from following Plato's command to 'know thyself', our minds are more oriented toward knowing about things in the world. My claim is that the fundamental world-directedness of mentality as a whole applies to sensory consciousness as well. Conscious sensory states represent the world at the present moment. Being cognitive creatures, our 'epistemic engines' will go to work on our conscious representations in their continuing efforts to make sense of things. Memory, language and first-person reports are instances of these cognitive labors. While these are important results of functioning conscious sensory states, they are not criterial for sensory consciousness, contrary to Dennett's claim.[143]

The beauty of Dennett's system is that it hangs together so perfectly. Start with first-person operationalism, an intuitively plausible approach. Who better than the subject of conscious sensory states to describe how they appear? Accept that first-person reports determine the truth about the content of sensory consciousness, and the distinction between Stalinesque and Orwellian forms of perceptual revision disappears. If a distinction can be drawn, it will certainly not be on the basis of first-person reports. Dennett and Kinsbourne are right that you can't tell the difference 'from the inside' (Dennett & Kinsbourne 1992:193). Then, without a distinction between Stalinesque and Orwellian forms of editing, there is nothing left but first-person reports to count as 'how things seem to you'. Subjectivity is reportability. What began as a methodological posit is now shown to be the case.

Take away first-person operationalism and the system falls apart. The only reason for believing first-person reports are always true is the methodological problem of determining reliable countervening evidence. This problem can be overcome, however, as it is in many cases where first-person reports are clearly false. First-person reports of memory can be false; people sometimes confabulate. First-person reports of current sensations can be false; people with Anton's syndrome deny their blindness. First-person reports of emotional state can be false; people can repress desires and feelings. Similarly, first-person reports of conscious sensory states are not in principle the last word on the contents of those states.

At the moment first-person reports are our primary resource in the investigation of sensory consciousness. Neuroscience, for all its bells and whistles, is not yet up to the task. But I am an optimist. I expect that in the forseeable future the current internecine squabbles among philosophers will result in an acceptable theory about exactly what conscious sensory states are (coordinated representations of the world at the present moment), and then the scientists can get on with the business of determining its neurological instantiation. If such a vision comes to pass, it may be possible to prove that the way things actually, objectively seem to you is not in fact the way they seem to seem to you. Stranger things have happened.

Appendix

A Speculative Hypothesis

One of the most exciting aspects of current research in consciousness theory is the interdisciplinary cross-pollination from a wide variety of fields. Advances in neuropsychology have been particularly influential in the philosophical investigation of consciousness. It is now considered a mark in favor of a philosophical theory of consciousness if it directly connects with empirical research. This appendix seeks to forge some of these connections. Inspired by Lycan's suggestion that we think of inner sensors as attention mechanisms (Lycan 1996: 14), I began to investigate psychological work on attention to see if any support for my theory of second sensing could be found. What I discovered is that 'attention' is used almost as broadly as 'consciousness', making the comparison of these two slippery topics quite challenging. Nonetheless, within the general category of 'attention theory' there is a great deal of research that suggests a connection between attention and consciousness. In what follows I draw on this research to speculate on a possible physical instantiation of conscious sensory states and the second sense mechanism that produces them. I begin by discussing the role of attention in stimulus integration and response preparation and compare this role to the proposed function of a second sense. I then use features of several neuropsychological theories of attention and consciousness to suggest how a particular sort of attention mechanism could perform the function of the second sense. I will not attempt to show that this is the best available neuropsychological description of either consciousness or attention. What I will argue is that second sensing and attention have important features in common, and these similarities provide some initial empirical support for the second sense theory of sensory consciousness.

A.1 Attention as the coordination of sensory representations

Although attentional phenomena are diverse and wide-ranging, three main functional categories form the heart of attention research: stimulus selection, response selection and sustained performance (Cohen 1993: 7–8; Cf. Parasura-

man 1998; Posner 1995) *Stimulus selection* highlights the role of attention in selecting and integrating stimuli. Given that the brain's processing capacity is limited, attention acts a filter mechanism or organizational device (Tstosos 2001; Reynolds & Desimone 2001; Niebur & Koch 1998; Treisman 1993). *Response selection* involves the role of attention in controlling available resources so as to prepare for an appropriate response. (Allport 1989, 1993; Tipper et al. 1999) *Sustained performance* is the ability to maintain attention on a particular target or task over time. (Parasuraman et al. 1998) In the following I will focus primarily on stimulus selection and integration. Response selection figures as a prominent factor in determining and facilitating selection, but the selection and integration of stimuli is of foremost interest in relation to the proposed function of the second sense.

One influential theory is Anne Treisman's Feature Integration Theory. (Treisman 1988, 1993, 1999; Treisman & Gelade 1980) On Treisman's account, simple features in a visual scene such as color, shape, orientation, and texture can be detected by parallel feature-recognition modules, but a conjunction of features (a blue square, for example) can only be detected by a serial process of focused attention. In early descriptions of the theory, Treisman suggested that parallel processing required no attentional resources (Treisman & Gelade 1980). More recently, she has determined that some form of attention – termed 'divided attention' – is required in order for parallel processes to form a coherent visual experience.

> We can never be aware of a free-floating orientation, colour, or shape. What varies across tasks is how broadly or narrowly focused the attention window is. Thus I assume . . . a continuum between divided and focused attention, as the size of the attention window narrows down or spreads wide. Texture segregation, visual pop-out, and detection of global alignment and shape . . . are done with a broad setting of the attention window, integrating feature maps at a global level. Accurate localization and conjoining of features for individual objects require narrowly focused attention. (Treisman 1993: 13–14)[144]

Treisman's distinction between divided and focused attention produces an analogous continuum to the one I describe between loosely coordinated and well-coordinated representations. By combining the two, we can describe the contrast in this way: When the attentional window is set widely, processing is rapid, integrating a broad range of salient features. Scene segmentation, surface orientation and pop-out features combine to form a loosely coordinated representation of the current environment. When the attentional win-

dow narrows, features identified at the attended location are combined into a well-coordinated representation that can serve to specify an object.

Because the relation between divided and focused attention is a continuum that varies according to task as well as over time, the contents of consciousness can shift instantaneously from well- to loosely- coordinated representations or vacillate among several more-or-less coordinated representations. Treisman has postulated the concept of an 'object file' to describe the way information about a particular object can influence attentional selection so as to produce more well-coordinated representations. Attending to an object establishes an object file, which is a representation of current properties at a location (Treisman 1993:24). An object file is then used to track changes in an object, as a template to aid conjunction detection for similar objects, or as a means of inhibiting response to irrelevant objects.

Ronald Rensink (2000) proposes a structure similar to object files, which he calls 'coherent objects.' On Rensink's account, low-level visual processes operate rapidly and in parallel to form 'proto-objects.' These are "relatively complex assemblies of fragments that correspond to localized structures in the world" (Rensink 2000:22). A great deal of detail is preserved at the level of proto-objects, but these representations are highly volatile, being instantly replaced when new stimuli appear on the retina (Rensink 2000:20). To form a more stable object representation – a 'coherent object' – requires focused attention to link proto-objects across space and time. Focused attention, guided by higher-level interpretive structures,[145] enables "a mapping between the ever-changing retinotopic coordinates of the proto-object and the more stable viewer- (or object-) centered coordinates of the [coherent object]" (Rensink 2000:24). According to Rensink, the relation between top-down interpretive influence and bottom-up visual processing is dynamic and interactive. The attentional link between the proto-object level and the coherent object level establishes a "two-way transmission of information between these structures" (Rensink 2000:24).

Counter-intuitively, however, Rensink argues that attention can only be focused on one object at a time, and as soon as focused attention is withdrawn the coherence of the attended object dissolves. Why, then, do our conscious states represent multiple, richly detailed objects rather than only one object at a time? Rensink's answer is that we maintain a 'virtual representation' of a scene. Rather than maintain a detailed representation of every object in a scene, we need only a sparse spatial representation to allow us to locate an object should it be of interest (Rensink 2000:28–33).

This description of virtual representation is a bit minimalist for my taste. Although some prefer an ascetic ontology when accounting for the representational contents of consciousness (Dennett 1991), I find it difficult to accept that *only* the few letters within my present focal range are coherently represented and the rest is a bare spatial map. Nor does this seem a bullet we need to bite. Interestingly, Rensink's comparison of virtual representation to the information access procedures of the World Wide Web (2000:28) is open to a more vivid portrayal of the contents available in a virtual representation. My personal computer only has the capacity to access a few sites of the vast quantities of information available on the web. Yet the ability to locate sites quickly and load required information makes it appear that all these vast quantities are simultaneously present in my own computer. It is impossible for my computer to store all of the information available, but there is no need to do so as long as it is easily located. This is Rensink's argument for minimalist virtual representation – why store it when you can acquire it just as quickly (Rensink 2000:28)? Note, though, that personal computers do store some information from sites accessed in a particular browsing session, even though only one or two sites are on screen at any time. I am well aware of this fact due to the frequent malfunctions of our limited capacity server. The message "Unable to contact server. Accessing previously cached version." appears when I attempt to return to a site after the server has crashed. If our virtual representations access the world in a way similar to the way personal computers access the web, perhaps they include several "previously cached" representations of available features, objects or scenes. Thus virtual representations could contain more information than simply scene structure yet less information than a fully detailed, constantly updated representation of every object.

In sum, we find several elements of attentional selection and integration correlate neatly with the coordination of representation into conscious states as described by the second sense theory. The distinction between focal and divided attention is comparable to the distinction between well- and loosely-coordinated representations. The volatile, momentary nature of proto-objects is consistent with the suggestion that conscious states represent the present moment, and are continually updated as new information becomes available. And the description of virtual representations as a sufficient representation to facilitate quick access to more detailed information is amenable to the proposal that conscious states represent more that what is available within foveal vision yet less than the full detail of an entire scene.

Though I have only considered two theories of attentional selection from a number of available accounts, the connections show the potential for explain-

ing aspects of sensory consciousness in terms of attention theory. This is an important result. Because neuropsychologists have concrete physiological theories about how attention works and what brain areas are involved, sensory consciousness loses some of its mystery through its association with attention. While there is not uniform agreement about which attention theory is correct, no one claims that attention is not physical or cannot be explained by objective methods of research. To the extent that sensory consciousness can be accounted for by theories of attention, it too escapes these theoretical mires.

A.2 Locating sensory consciousness

Since the second sense theory has provided a description of what sensory consciousness is – coordinated representations of the present moment – we might ask whether attention theory can provide any insight as to *where* sensory consciousness is. But first we have to determine how to look, since no single lesion or malfunction correlates consistently with loss of conscious sensory states. The standard way to determine the physical instantiation of a function is through dissociation tests. Researchers devise tests to see if damage to a particular physical region (an organ, muscle, etc.) causes a loss in function and conversely, if a loss in function can be traced to abnormalities in that same region. If damage to a physical region perfectly correlates with loss in function, researchers conclude that the function is instantiated in that region. Loss in sensory consciousness, however, does not perfectly correlate with damage to any particular physical region. To take conscious visual states as an example, deficits to the sensory mechanism – from retinal defects to lesions along the optic tract – result in complete blindness for some or all of the visual field (Kolb & Whishaw 1990). No information from the damaged part of the visual field is processed, either consciously or unconsciously.[146] Deficits to subcortical regions generate a variety of disfunctions. Lesions to the superior colliculus result in orienting deficits (Hilgetag et al. 2001; Johnson 1998; Zablocka & Zernicki 1996), while pulvinar damage causes difficulty in engaging attention and shifting attention (LaBerge 1995a, 2000). Thalamic lesions of the intralaminar nuclei can result in various sorts of confused states, coma and often death, depending on the extent of the damage (Bogen 1995).

Finally, deficits to cortical areas present the most perplexing malfunctions. Blindsight patients, who suffer from lesions to striate cortex, claim no awareness of any stimuli in the field affected by the lesion and will not voluntarily respond to stimuli presented there. If asked to guess, however, they are able to

correctly grasp or point to the presented object or choose its color or shape with a probability greater than chance (Weiskranz 1997; Farah 1997; Stoering & Cowey 1996). Patients with parietal damage also suffer from a partial deficit in consciousness, called hemineglect. Most common in cases of right parietal lesions, hemineglect patients claim no awareness of and do not respond spontaneously to stimuli on the contralesional side of space. A left neglect patient, for example, may bump into objects on the left, fail to read the left side of a page or eat food on the left side of the plate (Rafal 1998; Young & DeHaan 1993). Yet in a visual search task, patients with right parietal lesions were able to respond rapidly to distinct features on the neglected side of the display. It seems that subjects will continue to search a display if they know a target will appear, overcoming the attentional bias against the neglected side (Robertson 1998). Cases of apparently partial consciousness such as these are intriguing and worthy of consideration. Yet it is notoriously difficult to determine exactly how these deficits affect conscious states. Subject reports are all the evidence we have in these cases, and without a clear sense of how to interpret these reports we should be cautious about drawing conclusions based on them.[147]

We need not be so cautious in concluding that consciousness cannot be localized in a single brain area. In each of these cases sensory consciousness is clearly compromised in some way, yet each case involves a different part of the brain. Therefore, it is reasonable to conclude that there is no one part of the brain that is both necessary and sufficient for sensory consciousness.

A.2.A Many senses or one?

The lack of a single locus point for deficits of consciousness suggests there is no organ responsible for generating conscious states. Otherwise one would expect that deficits in consciousness could be traced to that organ. If no one organ generates conscious sensory states, then it seems there can be no 'inner' or 'second' sense that performs this function. Inner sense theorist David Armstrong has responded to this objection by comparing inner sensing to proprioception. No single organ is responsible for generating proprioceptive signals; nonetheless, proprioception is an important 'sixth' sense (Armstrong & Malcolm 1984:110). Lycan follows this response, proposing "ranks and ranks" of inner sensory mechanisms (Lycan 1987:72). The idea is to distribute the function of inner sensing broadly enough so that no single element need be identified as 'the' inner sense. In other words, on Lycan's approach we identify an inner sense by its function rather than by an end organ, and then we note that this function can be performed by many, many mechanisms.[148]

Certainly there is no physical reason why there could not be multiple inner sensory mechanisms busily producing representations of all sorts of different first-order states. But if there are really ranks and ranks of inner sensors, how does inner sensing perform its integrative function? Say that a different inner sense is responsible for producing representations of each different sort of state; minimally, we would have one sensor for visual sensations and another for auditory sensations. Then how is it possible to have a conscious sensation of looking at the lilac bush while hearing the clock chime? What makes these different inner sensory representations both part of the same conscious sensory state? Perhaps another inner sense combines the visual and auditory representations. If so, we would truly need ranks of inner sensors to accommodate all the possible combinations of sensations. More problematically, postulating a hierarchy of inner sensors all of which make sensory states conscious is inconsistent with the higher-order theorist claim that third-order sensings are different in kind than second-order sensings.[149]

More plausibly, sensory states could be integrated simply by being simultaneously scanned by some inner sense or another. In this case a conscious sensory state would be composed of all and only mental states being scanned by some inner sense at that moment. The addition of simultaneity is intriguing, but again it is important to consider how sensory states are integrated on this view. Neurological research shows that sensory signals are processed separately, each by its own dedicated system. The inner ear transduces auditory signals and sends this information to the temporal lobe (Wise et al. 2001; Kosmal 2000). The retina transduces visual signals which are then subdivided into two routes: the dorsal stream to cortical region V1 and the ventral stream to cortical region V4 (Humphreys 2001; Zeki 1993). Somewhere along the line (one might wonder exactly where or how) inner sensors must scan these states to form higher-order representations of them. At the moment of scanning, the sensory states become conscious and are thereby integrated with other states being simultaneously scanned.

One problem with the multiple sensor explanation is the difficulty in accounting for the selectivity of consciousness. We are not state conscious of everything in our environment at all times; only a fraction of available stimuli are represented by conscious states. Without some sort of selection process, our neurological system would be overloaded with information (Tsotsos et al. 2001). Further, if the function of sensory consciousness is to facilitate effective response, stimuli relevant to the task at hand need to be selected for conscious processing and distracting stimuli need to be inhibited. With multiple sensors, it is hard to see how this function could be performed. Some sort of conductor

mechanism would be needed to orchestrate which scanners operate when, otherwise there would be no way to limit which sensors are operating or to ensure relevant sensors are active. So functionally, if not anatomically, there must be a single second sense. An additional bonus of a functional description of the second sense is that it solves the question of how to look for sensory consciousness in the brain: we look for what mechanisms could do the job of second sensing. In the next section I propose an 'organ' of sorts to fulfill the selection function of the second sense, but this anatomical feature is not a necessary condition. The job of selection and integration could be distributed anatomically, so long as the distributed parts are functionally interconnected. By postulating a functional locus for second sensing, we avoid iterating levels of integration and control. A single second sense performs the integration and coordination functions necessary for sensory consciousness.

One final reminder is in order before we continue. Selection by the second sense is a necessary but not sufficient condition for sensory consciousness on the present account. In addition to the operation of a second sense, there must be the right sort of input (sensory states) and the right sort of output (coordinated representations of the world at the present moment). The joint sufficiency of these conditions explains why no single lesion produces all and only deficits in sensory consciousness, so multiple second sensors are not needed to answer this objection.

A.2.B Looking for the second sense

At this point we have enough of an idea about the function of the second sense in order to begin looking for a physical instantiation. Since, as the first section shows, second sensing shares some of the same functional characteristics as attention, we can look to neuropsychological theories of attention to see what might perform those functions.

Some of the most powerful models of the dynamic processes involved in attentional selection and integration involve an attention loop connecting midbrain and cortical structures (Posner 1994, 1995; Posner & DiGirolamo 2000; LaBerge 1995a, 1995b, 2000; Crick & Koch 1990a, 1990b, 1997; Llinás 1994, 1996, 2001). Attention is influenced both by bottom-up processes, where highly salient stimuli command attentional resources, and by top-down control, where attention is purposefully directed at a target (Miller 2000; LaBerge 2000). Thus, an adequate model of attention must account for both forms of influence, as well as the distinction between divided and focused attention discussed earlier. Loop, or circuit, theories meet these conditions by postulating

a dynamic relation among several brain areas (Reynolds & Desimone 2001; Edelman & Tononi 2001; Posner 1995b; LaBerge 1995a, 2000).

Michael Posner (1995b) has developed one of the most comprehensive circuit theories. Dividing attention into three functional categories – orientation, executive and alerting – Posner attributes these functions to the operation of three networks of anatomical areas (Posner 1995b:617). The orientation network is responsible for directing attention to targets, selecting stimuli and binding input according to spatial relations. According to Posner, this network, called the posterior attention network or PAN, is composed of the posterior parietal cortex, thalamic areas of pulvinar and reticular nuclei and parts of the superior colliculus (Posner & DiGirolamo 2000; Posner 1995b; Posner & Rothbart 1992). The element in this circuit relevant to second sensing is the subcortical modulation structure, specifically the modulatory function of the thalamus.[150]

Three features of thalamic structure and operation are suggestive of second sensing. First, the thalamus is widely connected to both sensory systems and cortex. Specific nuclei in the thalamus are dedicated to processing signals from each sensory system: the lateral geniculate nucleus (LGN) relays signals from the eyes; the medial geniculate nucleus (MGN) relays signals from the ear; the ventral posterior medial nucleus (VPM) relays gustatory signals. Almost every cortical area sends and receives signals from some thalamic nucleus (Newman 1996; Churchland 1995). The breadth and multiplicity of interconnections is unparalleled by any other brain structure (Bogen 1998), making the thalamus well suited to serve the global integration tasks required of the second sense.

Second, one section of the thalamus, the pulvinar, has been implicated in the selection and control of attention (Posner & DiGirolamo 2000). By enhancing or suppressing the inhibitory activity of the pulvinar, the speed of attentional shifts can be respectively decreased or increased. Monitoring of attentional shifts through positron emission tomography (PET) studies indicates that the pulvinar plays a role in selecting a stimulus from a crowded visual array (LaBerge 1995a, 2000; Colby 1991). One study (Olshausen et al. 1993) proposes the pulvinar as a seat of 'control neurons' designed to dynamically modify the connection strength of the neurons forming sensory pathways. In the visual system, when retinal images shift across saccades, the control neurons selectively modify intracortical synaptic strengths so as to produce an object representation that is stable both in position and scale. To put this description in more familiar terms, when I stand in the produce section and scan an apple for rotten sections, the pulvinar is important in forming the representation of a single apple from the various visual images of right, left, top and bottom. While

this is not the sort of integration function proposed for second sensing,[151] the role of the pulvinar in modifying synaptic strengths supports the idea that the thalamus is capable of and responsible for selecting and enhancing stimuli.

Yet neither of the information relay/modulation functions described above seems to require sensory consciousness. Surely, one might say, these processes could go on within me entirely without any awareness, just as digestion requires no awareness in order to perform its function. The descriptions support a connection between the thalamus and stimulus selection and control, but they do not show that mental states selected by the thalamus are conscious sensory states. To add to the intuitive plausibility that these are indeed conscious sensory states, consider the distinction between sleep and wakeful states. Deep, dreamless sleep has none of the markers of a conscious sensory state. The sleeper cannot remember (and therefore cannot report) her activity during this sort of sleep. Time appears to pass instantly. Some sensory responses occur, but they are limited to automatic responses such as changing position in order to ease a body cramp. Wakefulness, by contrast, carries all of the markers of conscious sensory states. Time is marked by events which can (to some extent) be reported immediately and, if sufficiently interesting, can be remembered later. Sensory responses are multi-modally coordinated, goal-directed actions. The sleep/wake distinction is clearly not equivalent to the distinction between unconscious and conscious states: dream states while sleeping might be conscious, and many unconscious states occur when one is awake. Nonetheless, the markers for dreamless sleep provide a useful contrast with the markers of wakefulness in illustrating some of the differences between unconscious and conscious states.

In this albeit impressionistic way the third role of the thalamus, regulating the sleep-wake cycle, adds intuitive force to the suggestion that the thalamus is involved in the production of conscious sensory states. Whereas wakefulness is characterized by coherent magnetic activity, deep sleep shows little such activity and instead exhibits delta wave patterns in EEG recordings (Llinás 2001; Llinás & Ribary 1993). Further, patterns of magnetic activity are reset after sensory input during wakeful states, but are not reset during either deep sleep or REM sleep (Llinás 2001; Llinás et al. 1994). These different forms of electrical activity are likely regulated by the thalamus. In sleep the thalamus disrupts sensory signals by generating rhythmic bursts of synchronized, low-frequency neuronal spikes. The rhythmic activity spreads throughout the cortex, blocking the relay of sensory information normally conducted through the thalamus. On waking, activity in the thalamus shifts to de-synchronized, high-frequency waves of low amplitude. Tonic firing, in contrast to burst firing during sleep,

allows the preservation of sensory input patterns (LaBerge 1995a: 180ff). The changes in rhythm of thalamic activity over the course of the sleep-wake cycle, from burst to tonic firing pattern, is another indicator that the thalamus may well perform the duties of the second sense.[152]

We should keep in mind that the whole neurobiological story about the sleep-wake cycle is more complicated than the changes to one organ. Given the interconnections, redundancies and plasticity of most brain processes, correlations between function and instantiation are rarely definitive. This is true for sensory consciousness as well as for the sleep/wake cycle. Therefore, other candidates for the role of second sense may supplant the thalamus on further study. For example, the reticular nucleus, a thin membrane that surrounds the thalamus, seems to function as an inhibitory mechanism in attention. It dampens signals that are deemed extraneous for one reason or another. The reticular nucleus also seems to play an important role in general arousal and vigilance (Knight & Grabowecky 2000; Posner 1995a; Posner & DiGirolamo 2000). Moreover, as noted earlier, the superior colliculus has been implicated in orienting attention (Hilgetag et al. 2001; Johnson 1998; Zablocka & Zernicki 1996). Either of these attentional mechanisms may eventually prove more suited to serve the function of second sensing, or all may function as a unit. Additionally, it is possible that aminergic neurons of the locus coeruleus are responsible for the tonic firing of the thalamus during wakefulness, and their lack of activity during sleep results in the synchronized burst firing that disrupts sensory processing (Churchland 1988). If so, the locus coeruleus may be responsible for coordinating the sleep-wake cycle which would diminish the intuitive appeal of the thalamus hypothesis.

Nonetheless, to date the thalamus is by far the strongest candidate to fill the role of second sense. What we need for a second sense is a device that takes sensory states as input and produces conscious sensory states, which are coordinated representations of the world at the present moment. The role of the thalamus in modulating sensory information in concert with the cortex makes it ideally suited to serve the function of a second sense. No other organ is as deeply interconnected with sensory and cortical systems, is central to the control of attention and regulates the sleep/wake cycle. Still, some would argue that the second sense is superfluous to the explanation of conscious sensory states. Even if a second sense controls attention and regulates the sleep/wake cycle, so the argument goes, these functions are not necessary to sensory consciousness. On an alternate account, conscious sensory states arise spontaneously out of the activity of cortical neural networks, so no second sense is needed.[153] Since attentional circuits invariably involve cortical projections, it would be worth-

while to take a closer look at the relation between conscious sensory states and the cortex according to the second sense theory.

A.2.C A home for conscious sensory states

With some variation in specific components, attentional circuits generally involve sensory, subcortical and cortical systems. David LaBerge (1995a, 2000), for example, proposes a circuit that nicely captures the processing I envision from sensory input to conscious states as output. According to LaBerge, sensory systems relay signals to the thalamus where information is modulated according to task requirements. These signals are then expressed in cortical areas, several of which may be activated simultaneously. Research indicates that the thalamus activates areas of cortex by modulating the neural discharge of projection cells (LaBerge 1995a; Llinás & Páre 1996; Llinás et al. 1994). Note that the thalamus and cortex form a circuit; activity in cortex also feeds back to the thalamus, so that cortical activity influences thalamic activity as well as the reverse. Further, thalamic operations activate patterns of cells in the cortex rather than single cells in keeping with the idea that cortical representations of your grandmother, for instance, may be widely distributed rather than isolated in specific 'grandmother' cells (LaBerge 1995a: 107).

To bring this all back to the second sense model, if the thalamus is the second sense, then the cortical expressions of thalamic operations are conscious sensory states. Therefore, conscious sensory states are physically instantiated in cortical activation patterns. This result is exciting because it connects the second sense theory with research on the relation of consciousness to cortical neural patterns. In recent discussions about possible neurological instantiations of consciousness, a theme of three elements recurs in various arrangements. These elements are: stimulus integration, synchronous neuronal oscillations, and attentional mechanisms. Variations showcase one or another of the elements and enlist the others for accompaniment. The following description combines aspects of several variations in order to provide a sense of the relation between cortex and consciousness that I have in mind.

We begin with stimulus integration and synchronous neural oscillations. In the interest of efficiency, processing of sensory signals is broken up into multiple separate pathways. Visual signals, as is now well established, are separated into at least four different systems: motion, form, color, and location. The brain's distributed architecture allows for speedy decomposition and identification of stimuli, but it also creates a problem for neurologists. Once sensory signals are broken down into their several parts, how does the brain manage

to put them back together again so as to form one multi-sensory conscious state? One currently popular suggestion among neurologists is that various signals are integrated by the synchronous firing patterns of neurons distributed throughout the cortex (Singer 2000; Llinás et al. 1994; Steriade et al. 1991; Crick & Koch 1990b). So, if I am looking at the lilac bush outside my window, the neurons in cortical region V4 that represent the purples, greens and browns of the bush are firing in synchrony with neurons in region V3 that represent the bushy shape. Neurons are connected in groups, or 'functional clusters', organized not by physical proximity but dynamic interactivity (Edelman & Tononi 2000: 145). Oscillatory activity at a frequency of 40–70 Hz could promote the frequency- and phase-locking necessary for synchrony, especially if the firing strength of oscillating neurons is modulated by a centralized feedback unit like the thalamus (Crick & Koch 1990b). An oscillatory pattern is important because one burst can predict the occurrence of the next with some probability. It is this predictability that could facilitate the coordination of remote cells into a single, synchronous pattern (Singer 1996: 118).

Synchronous neural firing could then account for the coordination of sensory representations. An attention mechanism (such as the thalamus) could coordinate sensory representations by selectively exciting some neurons and inhibiting others. Well-timed boosts and drags in signal strength can speed or slow firing patterns, thereby enhancing synchronization probabilities (Singer 1996, 2000; Edelman & Tononi 2000). Using the above example, when I attend to the lilac bush rather than the mock orange bush or some other object within my visual field, signals from the lilac are modulated so as to produce the synchronous firing of neurons representing the various aspects of the bush. Extraneous signals, such as those representing the mock orange or auditory signals from the computer and desk clock, are simultaneously disrupted to inhibit synchronization.

So, if conscious sensory states are coordinated representations, neural synchronization theory could explain how they are produced. An attention mechanism modulates sensory signals to promote synchronous firing patterns. The neurons bound in these patterns constitute conscious sensory states, and the features represented by those patterns form the content of the conscious states thereby instantiated. Sounds good. Unfortunately, though, this proposal has two potentially problematic implications.

1. A sensory state should be conscious when and only when there is a synchronous neuronal firing pattern produced by the second sense. Research cited earlier on the sleep-wake cycle supports this implication by showing the re-

lation of thalamic tonic firing with wakeful states, and burst firing with sleep states. As things stand, however, these studies are inconclusive. More evidence is needed to prove that the thalamus rather than some other mechanism is responsible for generating tonic and burst firing patterns. Moreover, we need to explain how the tonic patterns of wakeful states differ from the burst patterns of dreamless sleep states, since both sorts of pattern are highly synchronous. (Edelman & Tononi 2000). Several possibilities suggest themselves. It could be that a particular oscillation frequency, such as 40 Hz, is required for sensory consciousness, and burst firing oscillates at a different frequency. Or perhaps the type of synchronous pattern is critical. The burst pattern of deep sleep is global, all neurons fire synchronously, whereas the tonic pattern of wakefulness displays complex and variegated patterns of synchrony (Edelman & Tononi 2000). Clearly, we are still a long way off from a definitive description of the neuronal basis of sensory consciousness. Even so, the connections among neuronal patterns, attentional processes and conscious sensory states give us a place to start looking.

2. More worrisome, if neurons bound in synchronous firing patterns constitute conscious sensory states, then there seems to be no need for a second sense. A more simple theory of sensory consciousness would identify the appropriate firing patterns with conscious sensory states and leave it at that. On such a view, firing patterns synchronize spontaneously, by virtue of attractor fields or some such internal means of organization (Hardcastle 1996a, 1996b; Kinsbourne 1988, 1993, 1995). Being simpler, spontaneous formation of conscious sensory states seems the preferable solution. Nonetheless, there are several reasons to believe an additional mechanism is required. First, attractor fields do not explain why some sensory states are conscious and some are not. As Valerie Hardcastle describes the theory, incoming stimuli cause particular neurons to fire. The phase patterns of these neurons then serve as attractors, entraining other neurons to fire simultaneously in the same rhythm (Hardcastle 1996b: 258). But why are some stimuli successful attractors and others are not? Or successful at one moment and not at the next? Nor does the spontaneous generation of a particular frequency of oscillation, such as 40 Hz, supply a sufficient explanation of why some sensory states are conscious and others are not. Even if it is the case that all and only conscious sensory states exhibit a 40 Hz oscillation pattern, simply asserting the correlation is not explaining it. Some additional reason is needed to show why 40 Hz patterns are conscious and, more difficult for the spontaneous generation theorist, why some firing patterns oscillate at 40 Hz and some do not.

Attentional control is another reason that a second sense is necessary to conscious sensory states. Part of the phenomenon of sensory consciousness is our ability to direct its activity. At this moment I am state conscious of the words on the computer screen as I write this sentence, but in the next moment I may choose to turn my attention back to the lilac bush outside my window and so become state conscious of it. We have the ability to direct attention so as to be state conscious of some things rather than others. Some regulatory device is necessary to account for this ability.

One could object that attentional control is not necessary to have conscious sensory states, and so purposeful directedness should not be taken as part of the phenomena of sensory consciousness. On this view attention has the function of regulating sensory signals, focusing on some aspects of the sensory array and ignoring others. Conscious sensory states then spontaneously arise from the synchronous firing patterns of stimuli allowed to pass through the attentional gate. Conscious sensory states are completely separate from attentional processes, according to this hypothesis, because they occur after the work of attention is done. But this view arbitrarily separates two aspects of one functional process. One reason to think consciousness and attention are functionally connected is that the relation between sensory consciousness and attentional control is bi-directional. Not only does attentional control determine the content of conscious sensory states, but the reverse is true as well. Seeing the lilac bush might make me think of the lilacs in the vase which might remind me to get them some water. The chain of thoughts prompted by seeing the lilac bush directs my attention to the depletion of water in the vase.[154] The content of conscious sensory states influences attentional control even as attentional control selects the content of conscious sensory states.

Second, if conscious sensory states are representations of the present moment as I have claimed, their content depends on the selection and integration processes that are part and parcel of attentional control. On a functional account of representation, content is determined by what the representation does, and what it does depends on what things affect it and what it, in turn, affects. So, one can no more sever the functional relations between representations of the present moment and the selection processes that produce them than one can sever the functional relations between sensory representations and the transduction processes of the retina.[155] The interrelations of a functional system are what determine the nature of the parts, particularly in regard to representational content. While it is certainly appropriate to distinguish the functional role of the second sense from its product, sensory consciousness, this functional distinction in no way implies a functional isolation.

Sensory consciousness is distinct from the second sense, but the second sensing is nonetheless necessary for sensory consciousness.

In sum, the insufficiency of attractor fields and the importance of attentional control support the addition of a regulating device such as a second sense for selecting and integrating the content of conscious sensory states. Physiological evidence suggests the thalamus as a good choice for the role of second sense: the thalamus modulates sensory input, and its effects are expressed in the cortex, a happy home for conscious sensory states. The cortical firing patterns produced by these thalamic modulations instantiate conscious sensory states. We should not try to simplify the theory by saying that sensory consciousness is no more than cortical firing patterns, however. These patterns must be produced and regulated by the second sense in order to explain why some sensory states are conscious and some are not as well as how the content of conscious sensory states can be directed by (and direct) attention.

Thus concludes the speculative hypothesis. Attention, considered as input-coordinator oriented toward response, does exactly the work of second sensing. So we begin by identifying attention and second sensing. Within the context of circuit theories of attention we can then identify a likely candidate to serve as the second sense, the thalamus. Modulating input from sensory states, the thalamus expresses conscious sensory states across the cortex in the form of neural firing patterns. It is admittedly speculative, and perhaps astonishing, but possibly, just possibly, true.

Notes

1. Some may object to the notion that thoughts are felt in any way at all. No matter. This initial canvas of usage will be whittled and tailored to focus on a particular form of consciousness for examination.

2. I use 'intentional' and 'representational' interchangeably to mean 'of or about something'. I will assume that intentional states can be inherently unconscious, not even potentially conscious, contra Searle (1991, 1992).

3. Although I will use objects as examples of sensory representation, it may be that object representation requires higher cognitive abilities that the senses can muster, strictly speaking. If so, the theory will have a more narrow explanandum than my examples suggest.

4. In the most common description of the problem, the inverted spectrum is the possibility that two people could have identical functional systems yet experience different qualitative character. Whether or not an inverted spectrum is possible, the answer is to be found in a different kind of investigation than the one I am conducting. Whether someone could say they see red (and instantiate all of the functional roles of 'red') when in fact they experience green is a separate question from the question of whether they are experiencing at all. It is the latter question that concerns me here.

5. On another construal, reddish and greenish are essentially conscious and so are not so easily separated from the explanation of conscious sensory states. But such a view arbitrarily separates factors in common between unconscious and conscious sensory states. In the next section I will provide further reasons against identifying sensory qualities and conscious sensory states. For a fuller argument, see Rosenthal 1991a.

6. See Chapter 2 for more discussion of blindsight as evidence for unconscious sensory states.

7. In conversation.

8. Though I espouse a representationalist view, I mean to be neutral here. So if you object to the language of representation, you may substitute whatever form of non-representational language is preferable. For example, one could say 'when one has red or green visual sensation' or 'when one is seeing red or green.'

9. Dan Blair, in conversation.

10. Some may argue that there is 'nothing it is like' to have conscious thoughts and so this is a misleading characterization of an unconscious state. I believe there is something it is like to have a conscious thought, but this is not the place to argue the point. I have limited my discussion to conscious sensory states, and most agree there is 'something it is like' to have a conscious sensory state. On occasion I will discuss conscious and unconscious thoughts as

they were historically central to the debate about consciousness, but nothing in my argument hinges on whether or not there is something it is like to have a conscious thought.

11. Rosenthal translates this passage as: "no thought can exist in us of which we are not conscious at the very moment it exists in us." (Rosenthal 1997:747, Note 2)

12. The form of this result varies depending on the sort of test conducted. On some tests the subject will issue a report that there was no stimulus. On others the report will list other stimuli and leave out the sub-threshold stimulus. Or the subject might issue no report at all.

13. For a critique of this and several proposed empirical distinctions between unconscious and conscious processes, see Reingold & Merikle 1993.

14. Some philosophers have chosen this route. David Chalmers, for example, takes consciousness to be a primitive property on a par with the most basic physical properties. (Chalmers 1996) See Chapter 5 for more on this position.

15. 'Creature consciousness' is a property of creatures. It is, Rosenthal says, "the opposite of being asleep or knocked out." (Rosenthal 1993b:355)

16. See Chapter 2 for the full quote and an argument against the notion that thoughts 'obviously' involve some form of consciousness.

17. In order to continue using the handy grammatical form 'conscious of', I will later introduce the technical 'state conscious of' to refer exclusively to persons with conscious states that are of or about something. When I am 'conscious of' something, in this sense, I am in a conscious state about that thing.

18. 'Intransitive state consciousness' is the form of consciousness that requires explanation, in Rosenthal's taxonomy. I consider sensory consciousness to be a subset of Rosenthal's intransitive state consciousness. Conscious sensory states are a type of conscious state. Conscious thoughts are another type of conscious state, but one that I will not consider in this book.

19. See footnote 2.

20. In responding to commentary, Block (1997:159) revises this condition to 'voluntary or direct' control. Block wants to rule out blindsight as a case of A-consciousness but does not mean to suggest that states used in poor reasoning are not A-conscious.

21. As noted above, though, even the case of *actually* being used in reasoning or action may not make a state a conscious state. Repressed desires may guide my actions and unconscious thoughts may help solve a puzzle. Unconscious speech is less common, since a person is likely to become conscious of the states that generated the speech by hearing it.

22. 'External' here means 'external to the mind', where a mind is a system of mental representation. So, even a seemingly purely internal event, such as a wave of nausea, could represent an external object, viz. one's upset stomach. Michael Tye (1992, 1995, 1998, 2000) has provided an extended account of how apparently internal sensations are better understood as representations of external objects and events.

23. I am grateful to Crawford Elder for this example.

24. Of course, one could call such confabulations 'false memories' in which case all true episodic memories indicate conscious states at the time of the event remembered. This

would not be much help, however, since most episodic memories are distorted in some way. Moreover, it is impossible from a subjective point of view to determine which are the true and which are the false (parts of) memories.

25. We will return to the memory condition in Chapters 3 and 5 where we consider Dennett's discussion of the role of memory in determining the contents of conscious states.

26. Lycan also uses 'what it's like' to identify the target of his explanation. (Lycan 1996:4) More on the relation between Lycan's theory and my own later in this chapter.

27. Some people do not believe there is 'something it is like' to have a conscious thought. Rosenthal does not address this point, but the objection is another good reason to restrict ourselves to conscious sensory states.

28. Some qualification is in order here. Self-conscious or introspective states, to be discussed shortly, can also be conscious, and these states represent internal objects, namely mental states. This sort of state is rare and, like conscious thoughts, raises problems of its own. In the name of expedience, therefore, I will put them aside. Further, Rosenthal does not think that sensations are representational (Rosenthal 1993c: 202) and so they do not represent external objects on his view. Nonetheless it is sensory and representational states that become conscious according to Rosenthal's higher-order thought theory and representational states do represent external objects. A final concern is the possibility of conscious sensory states representing things that don't exist, internally or externally. As I have noted, one reason to adopt a teleofunctional semantics is its account of representational content in the absence of object represented. As far as I can see, though, any theory of intentional inexistents could apply here as well.

29. Again, 'the world' includes states of one's body. See Note 22.

30. These terms are translations from Herbert Davidson (1992:89).

31. The retentive imagination stores images from the *sensus communis*. The estimative faculty perceives something like Gibsonian *affordances* (see Chapter 4). Through this faculty, for example, the sheep perceives the harm of the wolf. (Avicenna 10th c/1952:30) Memory stores the perceptions of the estimative faculty, and the compositive imagination combines and divides forms from the other faculties. (Davidson 1992:89)

32. I toyed with the idea of calling this the 'common sense explanation of sensory consciousness.' But I haven't the style of Daniel Dennett to carry off such a bold title.

33. The distinction between sensation and cognition is a conceptual distinction rather than a practical one. Sensation seems to be cognitively penetrable, at least to some extent, so there is no sharp line between sensation and cognition. I do, however, believe that some non-conceptual sensation is possible and give some argument for this position in Chapter 5.

34. This last feature will be rejected by those who think that senses are fully non-representational. In such a case, I believe the first two features can bear the weight of the comparison.

35. One could also individuate an object purely conceptually, by learning its name, for example, that clearly does not involve perceptually isolating the borders of the thing. As I am looking at the distinction between conception and sensation here, only the requirements for sensation-based concepts are relevant to the point.

36. Imagery presents an interesting hybrid case as it represents an object in absence, yet is more detailed than conceptual representation. Here I would say that, though the physical object is absent, an intentional one is present. Thus the detail of the imagined object vanishes when the imagination ceases.

37. One could argue that mental states are self-intimating – simply by having them we are conscious of having them. But such a theory faces the difficult question of why some states intimate themselves and others don't, and why some intimate themselves intermittently.

38. Although Dretske focuses on conscious sensory states, what he calls 'perceptual experience', just as I do, (Dretske 1995:103) he frames his arguments in terms of 'state consciousness' generally. So I will use the more general term in describing his views as well as describing higher-order theories.

39. For higher-order theorists conscious states represent the world as well, but nothing about this relation is essential to their being conscious.

40. And shortly following: "S sees (hears, etc.) x (or that p) → S is conscious of x (that p)." (Dretske 1993b:265)

41. The key proponents in this debate about higher-order theory, Rosenthal (1991c, 1993b, 1997), Lycan (1996) and Dretske (1993a, 1993b, 1995), all frame their discussion in terms of state consciousness, so I will use this general term in describing their theories. As noted in Chapter 1, I take sensory consciousness to be a subset of state consciousness, and I will use this more specific term where applicable.

42. Originally coined by Malcolm (Armstrong & Malcolm 1984), these terms are used, with some variations in definition, by all the key parties: Armstrong (Armstrong & Malcolm 1984), Rosenthal (1991c, 1993b, 1997), Lycan (1996) and Dretske (1993a, 1993b, 1995).

43. See also Rosenthal 1997:741. At one point Rosenthal offered a minimalist form of the self-reference involved as "whatever individual has this very thought is also in the target mental state." (1991c:469) Later, however, Rosenthal disavowed this construal as problematically suggestive of self-intimational view of state consciousness. (Rosenthal 1993a:165)

44. "A mental state is non-introspectively conscious when accompanied by a relevant HOT; introspection occurs when there is a third-order thought that makes the second order thought conscious." (Rosenthal 1997:745)

45. Daniel Dennett (1991:317) has argued that this hierarchy is problematic because it allows for multiple sources of error. It seems, however, that multiple sources of error exist; there are various ways in which we are mistaken about our mental states. As I will argue in Chapter 5, Dennett's objection rests on his claim (which I dispute) that there is no more to state consciousness than reportability.

46. Rosenthal raises this objection against the dispositional higher-order thought theory. (Rosenthal 2002)

47. Carruthers accepts the standard developmental theory that children acquire the ability for higher-order thinking between the ages of 2 and 4 years of age. (Carruthers 2000:244–246)

48. Nonetheless, the dual content proposal is worth exploring in connection with the subjectivity of sensory states. As I argue in Chapter 4, sensory states carry information about

both subject and object by specifying object features relative to the sensing subject. The development of a higher-order ability to reflect on sensory content could explain how we are able to use the same sensory states either to represent the world – *red* – or to represent how we are representing – *as red*.

49. For brevity, I will use 'objects' or 'physical objects' as a shorthand for anything other than a mental state that can be the intentional object of a mental state, be it physical objects, events, states of affairs or whatever your ontology will allow.

50. Lycan is helpfully explicit on this point. After introducing use (3) 'Consciousness of' something, Lycan suggests that use (3) is not the problematic sense of 'consciousness' that many worry about. Rather it is just a special case of intentionality. (Lycan 1996:5)

51. Rosenthal comes even closer to identifying transitive consciousness and mental representation when he says, "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." (Rosenthal 1997:737) Still, though, there are conditions on transitive consciousness that are not standardly placed on mental representation, such as the stipulation that only a creature can be transitively conscious of something, mental states cannot be. (Rosenthal 1997:738)

52. I will refer only to the higher-order thought theory for the purpose of brevity, noting references to Lycan's higher-order perception theory in footnotes.

53. Lycan credits Dretske for this objection, which applies to both higher-order thought and higher-order perception theories. The higher-order perception description of the representation relation differs somewhat from the higher-order thought description and will be discussed in the next section. (Lycan 1996:23)

54. *Intransitive consciousness* in Rosenthal's lexicon is synonymous with state consciousness. (Rosenthal 1997:737) This sense is significantly removed from the original sense introduced by Malcolm as a conscious state with no object. (Armstrong & Malcolm 1984:30) But a conscious state could have an object, on Rosenthal's view, if it is a conscious thought about something. Of course Rosenthal is welcome to diverge from Malcolm's sense. The consequence of this terminology, however, seems to be that some (or all) cases of *in*transitive consciousness are also cases of *transitive* consciousness. Any conscious state that represents something is an instance of this oxymoronic category. For this and other reasons I will avoid the term 'intransitive consciousness' wherever possible.

55. Higher-order perception theorists have also committed themselves to unconscious mental states. See Armstrong 1984:127; Lycan 1996:16.

56. Lycan allows for some change in the mental state sensed as a result of higher-order perception (1996:35). However, there could not be significant change if it is to be the *same* mental state that is conscious by means of higher-order representation.

57. Rosenthal's more recent work on introspection and the self (2000, 2003) also suggests that this is where the mystery resides.

58. There is also a Higher-order Perception version of the view that conscious states are complex states composed of lower-order states and their higher-order representations. See Lormand, IS.

59. Alex Byrne poses a similar objection to higher-order thought theory. (1997:116)

60. For good reasons to reject the idea of a visual field other than the good old external world, see Austen Clark (2000)

61. I have been calling this the 'inner sense theory' in describing its relationship to the second sense theory I advocate. As I am focusing on higher-order structure in this chapter, I will use the term 'higher-order perception theory' throughout the chapter.

62. In the technical terms introduced above, perceptual consciousness involves first-order mental states without transitive consciousness of those states.

63. Introspection can be yet higher-order on either account, but only the lowest level of introspection is relevant to the present point.

64. Higher-order states also represent cognitive states and other higher-order states.

65. In Chapter 5 I tackle the general problem of 'zombies', hypothetical creatures who satisfy the conditions of the proposed explanation of sensory consciousness but who have no sensory consciousness.

66. In fact, the coordination function is essential to inner sensing. To constitute state consciousness, the output of a sensor must "contribute specifically to the integration of information in a way conducive to making the system's behavior appropriate to its input and circumstances." (Lycan 1996:32)

67. Güzeldere (1995), for example, suggests this as one interpretation of the higher-order perception theory.

68. These are roughly the objections to sense-data theory given in Barnes 1965.

69. What sort of isomorphism is necessary for representation is, of course, another hotly debated topic. But again, no specific answer is required to make the point. Lycan remains topic neutral, saying only that representations are 'normally caused by' the things they represent, where the normativity is cashed teleologically. (Lycan 1996:75)

70. Higher-order theorists might have such an explanation which I will consider in the next chapter as a possible objection to the second sense theory.

71. Thanks to Austen Clark for suggesting the term 'flat'.

72. "State consciousness implies that one is in some way or other transitively conscious of the state." (Rosenthal 1993b:359)

73. Dretske 1993b:265.

74. This is how Lycan has responded to Dretske's objection. (Lycan 1996:25ff)

75. For a persuasive argument against conflating sensation and state consciousness see Rosenthal 1991a.

76. Or even, perhaps, that it is a thing. As Dretske wonders, "If the concept one must have to be aware of something is a concept that applies to everything one can be aware of, what is the point of insisting that one must have it to be aware?" (Dretske 1993b:269)

77. See Dretske (1994, 1995: Chapter 4), as well as Neander's (1998) critique of this claim.

78. Not everyone agrees that Spot-sight presents a genuine phenomenon to be explained. Dennett himself discusses thimble-seeking precisely in order to show there are no Type 2 cases. I consider this objection in Chapter 5.

79. Cf. Dretske's Representational Thesis: All mental facts are representational facts. (Dretske 1995:xiii) and Lycan's 'hegemony of representation': "the mind has no special properties that are not exhausted by its representational properties, along with or in combination with the functional organization of its components." (Lycan 1996:11)

80. See Chapter 5 for Dennett's objection that Type 2 cases do not in fact exist.

81. As explained in Chapter 2, I use 'state conscious of *x*' only when one has a conscious state that is about *x*. When applicable, I use the term 'transitively conscious of *x*' to note the usage of 'conscious of' by higher-order theorists.

82. This is the second reading of Dretske, on the argument from Chapter 2.

83. I will make such an argument in Part 2.

84. See Chapter 5 for this argument.

85. See Chapter 2 for a fuller discussion of Rosenthal's proposal.

86. This is by no means intended as a theory of concepts, although I will say more about the basic distinction between sensory representation and conceptual representation later. I use the criterion of reidentification here to serve as a marker of the distinction between sensory representation and conceptual representation. Whatever else is involved in conceptual representation, it seems that at least this marker distinguishes it from sensory representation (Millikan 1998:59, 2000)

87. Whether it stays *the same* concept over time or not is another question.

88. Since I will claim that sensory consciousness arises in order to allow more sophisticated decision-making capacities, it is likely that conceptual development is coincident with the development of sensory consciousness ontogenetically. There is no necessary requirement, I claim, because there could be particular instances or aspects of sensory consciousness to which no concept is applied.

89. Thanks to Scott Sturgeon for anecdotal confirmation of this point.

90. The title recalls Stalin's show trials where false evidence was manufactured in order to convict political insurgents.

91. In Orwell's novel *1984* historians busily rewrote history to suit current governmental policy.

92. Several commentators to Dennett and Kinsbourne's article in *Behavioral and Brain Sciences* (1992) make this point. See especially Clark 1992:207–208; Glymour et al. 1992:209–210; Van Gulick 1992:228–229.

93. Dennett (1991:132) calls this the 'bizarre category of the objectively subjective' or 'how things seem to you even though they don't seem that way to you.' It is surprising that someone as anti-Cartesian as Dennett would espouse the discredited Cartesian notion of subjective infallibility, even in this limited context. More on this point in Chapter 5.

94. At this point one might wonder whether sensory consciousness is best described as a series of states or whether it is more like a single continuous process. The choice is not critical to the larger discussion, however. Like water in a stream, conscious sensory states flow continuously, periodically disrupted by sharp shifts in content and periods of unconsciousness. Nonetheless it is analytically useful to freeze the process so as to better examine the states

that compose it. For this analytical benefit I will speak of sensory consciousness in terms of individual states.

95. *How* the brain coordinates representations into conscious sensory states is a matter for the neurologists to determine, but see the Appendix for one speculative possibility.

96. See Baars et al. 1998; Baars 1993: 131; Baars and Fehling 1992: 204. Cf. Lahav 1993 for a similar account of the distinction between unconscious and conscious processing.

97. Lycan makes a similar connection in his suggestion that inner sensors are a type of attention mechanism. (Lycan 1996: 14). For more on the relation between sensory consciousness and attention, see the Appendix.

98. See Chapter 1 for an explanation of these three markers for sensation.

99. One common objection to calling the inner sense a 'sense' is that there is no obvious end organ. But, as Armstrong notes, there is no one end organ for proprioception. (Armstrong 1984: 111) Lycan makes the additional point that there is no claim here that inner sensing is like external senses in every single respect. (Lycan 1996: 28)

100. The scare quotes around 'consciousness' acknowledge Searle's (1980) classic argument against the sufficiency of syntax in determining semantics.

101. O'Brien and Opie (1999) offer a connectionist theory of consciousness that defends a distinction between unconscious and conscious states in terms of a difference in their vehicles. While I have reservations about some of their claims, their description of consciousness as a stable activation pattern in a neural net is consonant with the second sense theory.

102. See the Appendix for my argument for a single second sense as coordinating mechanism.

103. Dennett & Kinsbourne 1992: 194.

104. By this I mean that sensory stimuli from knee and foot generate conscious sensory states in the same brain. To what extent these relations form a self with the ability for self-reference is another question that I will not address.

105. 'Qualitative character' is a technical term designed to specify the increasingly vexed notion of 'qualia' mentioned in Chapter 1. As I use the term here, qualitative character is whatever features of a mental state determine how things appear. For instance, when something appears red rather than blue, this is a difference in the qualitative character of the relevant sensory states.

106. Though I follow the spirit of Wittgenstein here, I make no claim to be using particularly Wittgensteinian methods.

107. One possible exception is the question of privileged access: how is it that I know how things look to me but cannot know how things look to you? But as I will argue, this question also usually focuses on knowledge of qualitative character rather than the particular situation of the knower.

108. This is not to say I rule out a subjectivist account of qualitative character. Even an account that analyzes qualitative character in terms of features of the subject is still an objective description in the sense that the description, if not the qualitative character itself, is publicly accessible.

109. G. E. M. Anscombe (1974/1994) traces this mistake back to Descartes' *cogito* and the consequent fruitless search for the object to which 'I' infallibly refers. Shoemaker (1996) seems to hold a position similar to Descartes regarding introspective knowledge. In the usual case, Shoemaker claims, there is no role for "awareness of oneself as an object to play in explaining my introspective knowledge that I am hungry, angry, or alarmed. This comes out in the fact that there is no possibility here of a *mis*identification; if I have my usual access to my hunger, there is no room for the thought 'Someone is hungry all right, but is it me?'" (Shoemaker 1996:211)

110. I am assuming a direct reference theory of language here which I will not defend. (Kripke 1980; Putnam 1975) Readers who hold that words must be endowed with some kind of conscious speaker meaning to successfully refer may not be convinced by my distinction between grammatical and epistemic subjectivity. However, the main argument in this chapter does not stand or fall with this point.

111. Lycan (1996:71) has staunchly defended this position.

112. Note, though, that there is a difference between having a representation and knowing or reporting a representation. The next two sections will discuss the extent to which having a representation yields knowledge or reportability of the representation.

113. It is a matter of some debate exactly what the problem is. See Lycan (1996:71) and Clark (2000) for two views.

114. It may even be that there is no bodily state at all that I represent. In the case of phantom pain, for example, the victim represents her limb as in pain, but there is no limb.

115. I will take advantage of the familiarity with Nagel's argument to provide a brief and focused presentation of his claims. As a result I will overlook much that is both subtle and confused in the view.

116. Nagel refers to 'consciousness' and 'conscious mental states' generally, so in discussing his views I will do the same. When appropriate I will move to my more restricted subject, conscious sensory states.

117. Lycan 1987:77.

118. Or not. If qualitative character is essentially subjective, determinable by no set of objective stimuli consistent across all human beings, then such fine definitional tasks could prove impossible. Words would have no consistent reference for any two people. But as we will see, the limits of language in capturing qualitative detail is not the critical point.

119. It is worth noting one interpretation of Nagel's point that does not seem plausible: that bat mental states *qua* representations are inseparable from the bat's introspective point of view. As I suggested earlier, it is possible for introspection to inaccurately represent the content of one's mental states. Moreover, it is far more difficult to believe that a bat is capable of introspection than it is to believe in bat consciousness.

120. Unconscious sensory states, where there is nothing it is like to have them, would be another way of picking out an object. And if you are inclined to say there is nothing it is like to have a conceptual representation, then this would be yet another way of picking out an object.

121. Are we representing the same facts about the apple? No, but presumably we could by switching places and being attuned to different features.

122. Here I disagree with Lycan who says that the representations would be syntactically similar, but such a re-wiring would 'never happen in the real world.' (Lycan 1996: Note 17, 172) I think the pronomial aspect of subjective representations supports a stronger position.

123. The form of self-ascription need not be explicit. But the sensation must be taken as one's own in order to issue in the appropriate behavioral responses. In other words, if Oscar fails to represent the sensation as his own (rather than as Twin Oscar's sensation, which it also is), he will not be able to respond to it appropriately. I will say more on this point in the final part of the chapter.

124. Current theories emphasizing the dynamic interaction between subject and object continue to develop this Gibsonian insight. (Hurley 1998, 2001; O'Regan and Noë 2001; Noë 2001.)

125. I also need to see that the water will quench my thirst, but discussion of the motivational aspect of perception and action would take us beyond the scope of the present topic.

126. The idea of an egocentric map is revolutionizing work in artificial intelligence. Designers are now far more interested in how robots manage to act in an environment than they once were. See Clark 1997 for some recent developments.

127. Bermúdez (1998, 2001) makes a similar use of egocentric maps to argue for a primitive form of subjectivity that can serve as a foundation for a more sophisticated form of self-consciousness.

128. If you believe the bat to be a highly intelligent creature capable of introspecting, then think of another animal incapable of introspection yet capable of moving about, say a one-year old child.

129. The story I am telling here is only the barest sketch of what would be an enormously complex tale. The process of determining which environmental features are variant and invariant would require calculating the input from all perceptual systems, including a critical system which we have not yet discussed, proprioception.

130. In deference to the remarkable work of Armstrong and Lycan let me say that I believe many of their claims about the structure and function of a higher-order sense would remain. The main change would be that the theory would no longer purport to explain sensory consciousness. Sensory states are not conscious by being represented by higher-order states.

131. In this chapter I will often revert to the more general form 'consciousness' rather than my technical term 'sensory consciousness'. In some cases, consciousness theory generally is at issue, and in others it is unclear whether sensory consciousness or some other form is in question.

132. Alas, even here there is contest. Lycan has taken the idea of 'degrees' of consciousness to the extreme of saying that one internal monitor would count as a "little bit" of consciousness. (Lycan 1996: 40)

133. I have expressed my sympathies for a teleo-functional account of mental representation such as is offered by Millikan (1984, 1993), but zombie worries apply generally to representa-

tionalist theories of consciousness, so the response I offer here is a general representationalist reply.

134. Sophisticated computers, such as Deep Blue the chess master, present a more convincing case for possessing genuine representations. Yet even Deep Blue would not satisfy the second condition for sensory consciousness, to be discussed next.

135. As a functional theory, the second sense view is committed only to an identity between functional states (coordinated representations of the present moment, produced by a second sense) and mental states (conscious sensory states). In the Appendix I offer my own speculative ideas about a possible physical realization of these states. The point holds in either case, theoretical identities are justified by their explanatory power, and to date consciousness theories are relatively weak.

136. Rosenthal later drops the appeal to intuitions and unequivocably states the equivalence between consciousness and reportability. (Rosenthal 1997:747)

137. See Chapter 3.

138. As noted in Chapter 2, Dretske does not clearly distinguish between unconscious states and conscious states, so it is unclear whether 'awareness' necessarily involves a conscious state, as one might reasonably suppose.

139. See Chapter 2, Part 1, Section B.

140. According to Dennett, a subject must have "a certain family of concepts" to be conscious. (1994:550)

141. See Sellars (1963:141ff) for further development of this point.

142. Dennett and Kinsbourne (1992:234f) call them 'microtakings' or 'microjudgements,' content discriminations that have consequences for guiding action and forming a first-person narrative of memories, beliefs and reports.

143. Dennett is unequivocal: "'writing it down' in memory [is] criterial for consciousness . . . There is no reality of conscious experience independent of the effects of various vehicles of content on subsequent action (and hence, of course, on memory)." (Dennett 1991:132)

144. The metaphor of 'attention window' here is somewhat problematic as it implies a spatially contiguous region where only things 'inside' the window are attended. As Driver and Baylis (1998) argue, there is evidence that noncontiguous regions can be selected simultaneously. But even if scene segmentation involves a global level of pattern analysis (see Rensink's (2000) triadic architecture, described below), it is still reasonable to distinguish between scale (size of areas) and resolution (detail). (Cf. Nakayama & Joseph 1998:289). Thus the use of 'window' should be read as a metaphor, not as a literal description of the spatial relations involved.

145. Rensink describes a triadic architecture to account for the direction of attention in perception. The first, low-level system generates proto-objects; the second system forms proto-objects into coherent objects, and the third system provides a setting to guide attention. Rensink considers the setting system to be non-attentional, relying on higher-level interpretive structures such as concepts and memories to determine salient features of context (Rensink 2000:34–37). Note that the levels here refer to levels of processing, and should

not be confused with the orders of representation proposed by higher-order theories of consciousness.

146. But retinal blindness does not mean that conscious visual states are impossible. Direct stimulation of the visual cortex produces a similar experience in retinally blind subjects as in fully sighted subjects. (Cowey & Walsh 2000)

147. For example, Stoering and Barth (2001) found that patients will report that they are 'aware of' stimuli in the blind field, but will withhold the claim that they 'see' stimuli presented in the same location.

148. Additionally, each mechanism could be composed of widely distributed parts. Visual representation is not restricted to retinal representations, for example. The individuation is functional, not anatomical.

149. Lycan does not commit himself to this claim, but Armstrong does (Armstrong & Malcolm 1984: 121) as does higher-order thought theorist David Rosenthal (1997: 730, 745f).

150. Two other important subcortical structures – reticular nucleus and superior colliculus – will be considered shortly.

151. Several different sorts of integration have been proposed as functions of consciousness: binding of multiple features into a single, stable object over time (the sort of binding described here); binding of current thoughts and experiences with past and future thoughts/experiences (Myin 1998); binding of multiple features into a coordinated representation of the present moment. I propose the last form of binding as the principal function of sensory consciousness, although it is not unrelated to the first form.

152. However, for evidence against the role of the thalamus in producing the activity characterizing wakeful states, see Gray & Singer 1989.

153. Though not arguing specifically against a second sense, several theorists have suggested spontaneous neural networks or neural fields as an explanation of consciousness. See, for example, Kinsbourne 1988, 1993, 1995; Hardcastle 1996a; Greenfield 1995, 1997. I argue explicitly against this suggestion in the next section.

154. All of this could happen unconsciously as well. But if such a train of thought can be conscious, which seems it can, then conscious states can sometimes direct attention.

155. Nor can we separate sensory or conscious representation from the effects these representations have on the rest of the mental system.

References

Allport, Alan (1989). Visual attention. In M. Posner (Ed.), *Foundations of cognitive science* (pp. 632–681). Cambridge, MA: MIT Press.

Allport, Alan (1993). Attention and control: Have we been asking the wrong questions? A critical review of twenty-five years. In D. E. Meyer, & S. Kornblum (Eds.), *Attention and performance XIV: Synergies in experimental psychology, artifical intelligence, and cognitive science* (pp. 183–218). Cambridge, MA: MIT Press.

Anscombe, G. E. M. (1974/1994). *The first person*. Ed. Quassim Cassam. Oxford: Oxford University Press.

Aquinas, St. Thomas (1273/1989). *Summa Theologicae: A concise translation*. Ed. and Trans. Timothy McDermott. Westminster, MD: Christian Classics.

Aristotle (4th c. BC/1973). *Introduction to Aristotle*. Ed. Richard McKeon. Chicago: University of Chicago Press.

Armstrong, David M. (1981). *The nature of mind and other essays*. Ithaca: Cornell University Press.

Armstrong, David M. (1968/1993). *A materialist theory of the mind*. London: Routledge.

Armstrong, David M. (Unpublished manuscript, TW). The three waves.

Armstrong, David M., & Norman Malcolm (1984). *Consciousness and causality*. Oxford: Basil Blackwell.

Astington, Janet W., Paul L. Harris, & David R. Olson (Eds.). (1988). *Developing theories of mind*. Cambridge: Cambridge University Press.

Avicenna (10th c./1952). *Avicenna's psychology*. Trans. F. Rahman. London: Oxford University Press.

Baars, Bernard J. (1988). *A cognitive theory of consciousness*. Cambridge: Cambridge University Press.

Baars, Bernard J. (1993). How does a serial, integrated and very limited stream of consciousness emerge from a nervous system that is mostly unconscious, distributed, parallel and of enormous capacity? In G. Bock, & J. Marsh (Ed.), *Experimental and theoretical studies in consciousness* (pp. 282–303). Chichester: John Wiley and Sons.

Baars, Bernard J. (1997). *In the theater of consciousness: The workspace of the mind*. Oxford: Oxford University Press.

Baars, Bernard J., & Michael Fehling (1992). Consciousness is associated with central *as well as* distributed processes. *Behavioral and Brain Sciences, 15* (2), 203–204.

Baars, Bernard J., James Newman, & John G. Taylor (1998). Neuronal mechanisms of consciousness: A relational global workspace framework. In S. R. Hameroff, A. W. Kaszniak, & A. C. Scott (Eds.), *Toward a science of consciousness II: The second discussions and debates* (pp. 269–278). Cambridge, MA: MIT Press.

Baddeley, Alan, & Lawrence Weiskrantz (Eds.). (1993). *Attention: Selection, awareness and control*. Oxford: Clarendon Press.

Barnes, Winston (1965). The myth of sense data. In R. J. Swartz (Ed.), *Perceiving, sensing, and knowing* (pp. 138–167). New York: Doubleday and Co.

Bermúdez, José Luis (1998). *The paradox of self-consciousness*. Cambridge, MA: MIT Press.

Bermúdez, José Luis (2001). Nonconceptual self-consciousness and cognitive science. *Synthese, 129*, 129–149.

Block, Ned (1978/1991). Troubles with functionalism. In D. Rosenthal (Ed.), *Nature of Mind* (pp. 211–228). Oxford: Oxford University Press.

Block, Ned (1995). On a confusion about a function of consciousness. *Behavioral and Brain Sciences, 18*, 227–287.

Block, Ned (1997). Biology versus computation in the study of consciousness. *Behavioral and Brain Sciences, 20* (1), 159–166.

Block, Ned, Owen Flanagan, & Güven Güzeldere (Eds.). (1997). *The nature of consciousness: Philosophical debates*. Cambridge, MA: MIT Press.

Block, Ned, & Robert Stalnaker (1999). Conceptual analysis, dualism, and the explanatory gap. *Philosophical Review, 108* (1), 1–46.

Bock, Gregory, & Joan Marsh (Eds.). (1993). *Experimental and theoretical studies in consciousness*. Ciba Foundation Symposium, 174. Chichester: John Wiley and Sons.

Bogen, Joseph E. (1995). On the neurophysiology of consciousness: I. An overview. *Consciousness and Cognition, 4*, 52–62.

Bogen, Joseph E. (1998). Locating the subjectivity pump: The thalamic intralaminar nuclei. In S. R. Hameroff, A. W. Kaszniak, & A. C. Scott (Eds.), *Toward a science of consciousness II: The second discussions and debates* (pp. 237–246). Cambridge, MA: MIT Press.

Braun, Jochen, Christof Koch, & Joel L. Davis (Eds.). (2001). *Visual attention and cortical circuits*. Cambridge, MA: MIT Press.

Bridgeman, Bruce (1992). What is consciousness for, anyway? *Behavioral and Brain Sciences, 15* (2), 206–207.

Broad, C. D. (1923/1969). *Scientific thought*. New York: Humanities Press.

Buzsáke, G., R. Llinás, W. Singer, A. Berthoz, & Y. Christen (Eds.). (1994). *Temporal Coding in the Brain*. Berlin: Springer-Verlag.

Byrne, Alex (1997). Some like it HOT: Consciousness and higher-order thoughts. *Philosophical Studies, 86*, 103–129.

Campbell, John (1994). *Past, space and self*. Cambridge, MA: MIT Press.

Carruthers, Peter (2000). *Phenomenal consciousness: A naturalistic theory*. Cambridge: Cambridge University Press.

Chalmers, David J. (1996). *The conscious mind*. Oxford: Oxford University Press.

Chalmers, David J. (1997a). Facing up to the problem of consciousness. In J. Shear (Ed.), *Explaining consciousness: The "hard problem"* (pp. 9–32). Cambridge, MA: MIT Press.

Chalmers, David J. (1997b). Moving forward on the problem of consciousness. In J. Shear (Ed.), *Explaining consciousness: The "hard problem"* (pp. 379–422). Cambridge, MA: MIT Press.

Chalmers, David J., & Frank Jackson (2001). Conceptual analysis and reductive explanation. *Philosophical Review, 110* (3), 315–360.

Churchland, Patricia Smith (1988). Reduction and the neurobiological basis of consciousness. In A. J. Marcel, & E. Bisiach (Eds.), *Consciousness in contemporary science* (pp. 273–304). Oxford: Oxford University Press.

Churchland, Paul M. (1979). *Scientific realism and the plasticity of mind.* Cambridge: Cambridge University Press.

Churchland, Paul M. (1995). *The engine of reason, the seat of the soul.* Cambridge, MA: MIT Press.

Clark, Andy (1992). Experiential facts? *Behavioral and Brain Sciences, 15* (2), 207–208.

Clark, Andy (1997). *Being there: Putting brain, body and world together again.* Cambridge, MA: MIT Press.

Clark, Austen (1993). *Sensory qualities.* Oxford: Clarendon Press.

Clark, Austen (2000). *A theory of sentience.* Oxford: Oxford University Press.

Cohen, Ronald A. (1993). *The neuropsychology of attention.* New York: Plenum Press.

Colby, C. L. (1991). The neuroanatomy and neurophysiology of attention. *Journal of Child Neurology, 6*, 90–118.

Cowey, Alan, & Vincent Walsh (2000). Magnetically induced phosphenes in sighted, blind and blindsighted observers. *Neuroreport, 11* (14), 3269–3273.

Crane, Timothy (Ed.). (1992). *The contents of experience: Essays on perception.* Cambridge: Cambridge University Press.

Crick, Francis, & Christof Koch (1990a). Some reflections on visual awareness. *Cold Spring Harbor Symposia on Quantitative Biology, 55*, 953–962.

Crick, Francis, & Christof Koch (1990b). Towards a neurobiological theory of consciousness. *Seminars in the Neurosciences, 2* (4), 263–275.

Davidson, Herbert A. (1992). *Alfarabi, Avicenna and Averroes, on intellect.* Oxford: Oxford University Press.

Davies, Martin, & Glyn W. Humphreys (Eds.). (1993). *Consciousness: Psychological and philosophical essays.* Readings in Mind and Language, 2. Oxford: Basil Blackwell.

Dennett, Daniel, & Marcel Kinsbourne (1992). Time and the observer: The where and when of consciousness in the brain. *Behavioral and Brain Sciences, 15* (2), 183–247.

Dennett, Daniel C. (1991). *Consciousness explained.* Boston: Little, Brown and Company.

Dennett, Daniel C. (1994). Get real. *Philosophical Topics, 22* (1 & 2), 505–568.

Descartes, René (1641/1984). Author's replies to the fourth set of objections. In J. Cottingham, R. Stoothoff, & D. Murdoch (Trans.), *The philosophical writings of Descartes.* Cambridge: Cambridge University Press.

Dretske, Fred I. (1993a). The nature of thought. *Philosophical Studies, 70*, 185–199.

Dretske, Fred I. (1993b). Conscious experience. *Mind, 102* (406), 265–281.

Dretske, Fred I. (1994). Differences that make no difference. *Philosophical Topics, 22* (1–2), 41–57.

Dretske, Fred I. (1995). *Naturalizing the mind.* Cambridge, MA: MIT Press.

Dretske, Fred I. (1997). What good is consciousness? *Canadian Journal of Philosophy, 27* (1), 1–15.

Driver, Jon, & Gordon C. Baylis (1998). Attention and visual object segmentation. In R. Parasuraman (Ed.), *The attentive brain* (pp. 299–326). Cambridge, MA: MIT Press.

Edelman, Gerald M., & Giulio Tononi (2000). Reentry and the dynamic core: Neural correlates of conscious experience. In T. Metzinger (Ed.), *Neural correlates of consciousness* (pp. 139–151). Cambridge, MA: MIT Press.

Evans, Gareth (1982). The varieties of reference. J. McDowell (Ed.). Oxford: Clarendon Press.

Farah, Martha J. (1997). Visual perception and visual awareness after brain damage: A tutorial overview. In N. Block, O. Flanagan, & G. Güzuldere (Eds.), *The nature of consciousness: Philosophical debates* (pp. 203–236). Cambridge, MA: MIT Press.

Gazzaniga, Michael S. (Ed.). (1995). *Cognitive neurosciences*. Cambridge, MA: MIT Press.

Gazzaniga, Michael S. (Ed.). (2000). *The new cognitive neurosciences*. 2nd Edition. Cambridge, MA: MIT Press.

Gennaro, Rocco J. (1996). *Consciousness and self-consciousness: A defense of the higher-order thought theory of consciousness*. Amsterdam: John Benjamins Publishing.

Gibson, James J. (1986). *The ecological approach to visual perception*. Hillsdale, NJ: Lawrence Erlbaum Assoc.

Glymour, Bruce, Rick Grush, Valerie Gray Hardcastle, Brian Keeley, Joe Ramsey, Oron Shagrir, & Ellen Watson (1992). The Cartesian Theater stance. *Behavioral and Brain Sciences, 15* (2), 209–210.

Gray, Charles M., & Wolf Singer (1989). Stimulus-specific neuronal oscillations in orientation columns of cat visual cortex. *Proceedings of the Proceedings of the National Academy of Sciences of the United States of America, 86* (5), 1698–1702.

Greenfield, Susan (1995). *Journey to the centers of the mind: Toward a science of consciousness*. New York: W.H. Freeman and Co.

Greenfield, Susan (1997). How might the brain generate consciousness? *Communication and Cognition, 30* (3/4), 285–300.

Güzeldere, Güven (1995). Is consciousness the perception of what passes in one's own mind? In T. Metzinger (Ed.), *Conscious experience* (pp. 335–358). Paderborn: Ferdinand Schoeningh.

Hameroff, Stuart R., Alfred W. Kaszniak, & Alwyn C. Scott (Eds.). (1996). *Toward a science of consciousness: The first Tucson discussions and debates*. Cambridge, MA: MIT Press.

Hameroff, Stuart R., Alfred W. Kaszniak, & Alwyn C. Scott (Eds.). (1998). *Toward a science of consciousness II: The second discussions and debates*. Cambridge, MA: MIT Press.

Hardcastle, Valerie Gray (1996a). The binding problem and neurobiological oscillations. In S. R. Hameroff, A. W. Kaszniak, & A. C. Scott (Eds.), *Toward a science of consciousness: The first Tucson discussions and debates* (pp. 51–65). Cambridge, MA: MIT Press.

Hardcastle, Valerie Gray (1996b). How we get there from here: Dissolution of the binding problem. *Journal of Mind and Behavior, 17* (3), 251–266.

Hilgetag, Claus C., Stephen G. Lomber, & Bertram R. Payne (2001). Neural mechanisms of spatial attention in the cat. *Neurocomputing, 38–40*, 1281–1287.

Humphreys, Glyn W. (2001). A multi-stage account of binding in vision: Neuropsychological evidence. *Visual Cognition, 8* (3–5), 381–410.

Humphreys, Glyn W., John Duncan, & Anne Treisman (Eds.). (1999). *Attention, space and action: Studies in cognitive neuroscience*. Oxford: Oxford University Press.

Hurley, Susan L. (1998). *Consciousness in action*. Cambridge, MA: Harvard University Press.

Hurley, Susan L. (2001). Perception and action: Alternate views. *Synthese, 129*, 3–40.

Johnson, Mark H. (1998). Developing an attentive brain. In R. Parasuraman (Ed.), *The attentive brain* (pp. 427–444). Cambridge, MA: MIT Press.

Kant, Immanuel (1781/1990). *The critique of pure reason*. Trans. J. M. D. Meiklejohn. Buffalo: Prometheus.

Kinsbourne, Marcel (1988). Integrated field theory of consciousness. In A. J. Marcel, & E. Bisiach (Eds.), *Consciousness in contemporary science* (pp. 239–256). Oxford: Oxford University Press.

Kinsbourne, Marcel (1993). Integrated cortical field model of consciousness. In G. Bock, & J. Marsh (Eds.), *Experimental and theoretical studies in consciousness* (pp. 43–60). Chichester: John Wiley and Sons.

Kinsbourne, Marcel (1995). Models of consciousness: Serial or parallel in the brain? In M. S. Gazzaniga (Ed.), *Cognitive neurosciences* (pp. 1321–1329). Cambridge, MA: MIT Press.

Knight, Robert T., & Marcia Grabowecky (2000). Prefrontal cortex, time, and consciousness. In M. Gazzaniga (Ed.), *The new cognitive neurosciences, 2nd Edition* (pp. 1319–1339). Cambridge, MA: MIT Press.

Kolb, Bryan, & Ian Q. Whishaw (1990). *Fundamentals of human neuropsychology*, 3rd Edition. New York: W. H. Freeman and Co.

Kosmal, Anna (2000). Organization of connections underlying the processing of auditory information in the dog. *Progress in Neuro-Psychopharmacology and Biological Psychiatry, 24* (5), 825–854.

Kripke, Saul (1980). *Naming and necessity*. Oxford: Blackwell Publ.

LaBerge, David (1995a). *Attentional processing: The brain's art of mindfulness*. Cambridge, MA: Harvard University Press.

LaBerge, David (1995b). Computational and anatomical models of selective attention in object identification. In M. S. Gazzaniga (Ed.), *Cognitive neurosciences* (pp. 649–662). Cambridge, MA: MIT Press.

LaBerge, David (2000). Networks of attention. In M. Gazzaniga (Ed.), *The new cognitive neurosciences*, 2nd Edition. (pp. 711–725). Cambridge, MA: MIT Press.

Lahav, Ran (1993). What neuropsychology tells us about consciousness. *Philosophy of Science, 60*, 67–85.

Levine, Joseph (1997). Recent work on consciousness. *American Philosophical Quarterly, 34* (4), 379–404.

Llinás, Rodolfo (2001). *I of the vortex*. Cambridge, MA: MIT Press.

Llinás, Rodolfo, & Patricia Smith Churchland (Eds.). (1996). *The mind-brain continuum*. Cambridge, MA: MIT Press.

Llinás, Rodolfo, & D. Paré (1996). The brain as a closed system modulated by the senses. In R. Llinás, & P. S. Churchland (Eds.), *The mind-brain continuum* (pp. 1–18). Cambridge, MA: MIT Press.

Llinás, Rodolfo, & U. Ribary (1993). Coherent 40-Hz oscillation characterizes dream state in humans. *Proceedings of the National Academy of Sciences of the United States of America, 90*, 2078–2081.

Llinás, Rodolfo, U. Ribary, M. Joliot, & X.-J. Wang (1994). Content and context in temporal thalamocortical binding. In G. Buzsáke, R. Llinás, W. Singer, A. Berthoz, & Y. Christen (Eds.), *Temporal Coding in the Brain* (pp. 251–272). Berlin: Springer-Verlag.

Locke, John (1689/1975). *An essay concerning human understanding*. Ed. Peter Nidditch. Oxford: Clarenden Press.

Lormand, Eric (Unpublished manuscript, IS). Inner sense until proven guilty. http://www-personal.umich.edu/~lormand/phil/cons/inner_sense.htm

Lycan, William (1987). *Consciousness*. Cambridge, MA: MIT Press.

Lycan, William (1996). *Consciousness and experience*. Cambridge, MA: MIT Press.

Marcel, Anthony J. (1983). Conscious and unconscious perception: Experiments on visual masking and word recognition. *Cognitive Psychology, 15*, 197–237.

Marcel, Anthony J., & Edoardo Bisiach (Eds.). (1988). *Consciousness in contemporary science*. Oxford: Oxford University Press.

Maunsell, John H. R. (1995). The brain's visual world: Representation of visual targets in cerebral cortex. *Science, 270* (5237), 764–769.

Metzinger, Thomas (Ed.). (1995). *Conscious experience*. Paderborn: Ferdinand Schoeningh.

Metzinger, Thomas (Ed.). (2000). *Neural correlates of consciousness*. Cambridge, MA: MIT Press.

Meyer, David E., & Sylvan Kornblum (Eds.). (1993). *Attention and performance XIV: Synergies in experimental psychology, artifical intelligence, and cognitive science*. Cambridge, MA: MIT Press.

Miller, Earl K. (2000). The prefrontal cortex and cognitive control. *Nature Reviews, 1*, 59–65.

Millikan, Ruth Garrett (1984). *Language, thought, and other biological categories*. Cambridge, MA: MIT Press.

Millikan, Ruth Garrett (1993). *White queen psychology and other essays for Alice*. Cambridge, MA: MIT Press.

Millikan, Ruth Garrett (1998). A common structure for concepts of individuals, stuffs and real kinds: More Mama, more milk, and more mouse. *Behavioral and Brain Sciences, 21* (1), 55–100.

Millikan, Ruth Garrett (2000). *On clear and confused ideas*. Cambridge: Cambridge University Press.

Milner, A. D., & M. D. Rugg (Ed.). (1992). *The neuropsychology of consciousness*. London: Academic Press.

Milner, D., & M. Goodale (1995). *The visual brain in action*. Oxford: Oxford University Press.

Myin, Erik (1998). Holism, functionalism and visual awareness. *Communication and Cognition, 31* (1), 3–20.

Nagel, Thomas (1979). *Mortal questions*. Cambridge: Cambridge University Press.

Nagel, Thomas (1974/1991). What is it like to be a bat? In D. M. Rosenthal (Ed.), *The nature of mind* (pp. 422–428). Oxford: Oxford University Press.

Neander, Karen (1998). The division of phenomenal labor: A problem for representational theories of consciousness. In J. E. Tomberlin (Ed.), *Language, mind, and ontology* (pp. 411–434). Malden, MA: Blackwell Publishers.

Neisser, Ulric (Ed.). (1993). *The perceived self: Ecological and interpersonal sources of self knowledge*. Emory Symposia on Cognition, 5. Cambridge: Cambridge University Press.

Newman, James. (1996). Thalamocortical foundations of experience. Web page, [accessed April 2001]. Available at http://www.phil.vt.edu/assc/newman/.

Niebur, Ernst, & Christof Koch (1998). Computational architectures for attention. In R. Parasuraman (Ed.), *The attentive brain* (pp. 163–186). Cambridge, MA: MIT Press.

Noë, Alva (2001). Experience and the active mind. *Synthese, 129*, 41–60.

O'Brien, Gerard, & Jonathan Opie (1999). A connectionist theory of phenomenal experience. *Behavioral and Brain Sciences, 22*, 127–196.

O'Regan, J. Kevin, & Alva Noë (2001). What it is like to see: A sensorimotor theory of perceptual experience. *Synthese, 129*, 79–103.

Oshausen, Bruno A., Charles H. Anderson, & David C. Van Essen (1993). A neurobiological model of visual attention and invariant pattern recognition based on dynamic routing of information. *Journal of Neuroscience, 13* (11), 4700–4719.

Parasuraman, Raja (Ed.). (1998). *The attentive brain*. Cambridge, MA: MIT Press.

Parasuraman, Raja, Joel S. Warm, & Judi E. See (1998). Brain systems of vigilance. In R. Parasuraman (Ed.), *The attentive brain* (pp. 221–256). Cambridge, MA: MIT Press.

Posner, Michael I. (1995a). Attention in cognitive neuroscience: An overview. In M. S. Gazzaniga (Ed.), *Cognitive neurosciences* (pp. 615–624). Cambridge, MA: MIT Press.

Posner, Michael I. (1995b). Introduction. In M. S. Gazzaniga (Ed.), *Cognitive neurosciences* (pp. 613–614). Cambridge, MA: MIT Press.

Posner, Michael I., & Gregory DiGirolamo (2000). Attention and cognitive neuroscience: An overview. In M. Gazzaniga (Ed.), *The new cognitive neurosciences*, 2nd Edition (pp. 623–631). Cambridge, MA: MIT Press.

Posner, Michael I., & Marcus E. Raichle (1994). *Images of mind*. New York: Scientific American Library.

Posner, Michael I., & Mary K. Rothbart (1992). Attentional mechanisms and conscious experience. In A. D. Milner, & M. D. Rugg (Eds.), *The neuropsychology of consciousness* (pp. 91–111). London: Academic Press.

Putnam, Hilary (1975). The meaning of 'meaning'. In K. Gunderson (Ed.), *Language, Mind and Knowledge: Minnesota Studies in Philosophy of Science*, Vol. 7. Minneapolis: University of Minnesota Press.

Rafal, Robert D. (1998). Neglect. In R. Parasuraman (Ed.), *The attentive brain* (pp. 489–526). Cambridge, MA: MIT Press.

Rahman, F. (1952). *Avicenna's psychology*. London: Oxford University Press.

Reingold, Eyal M., & Philip M. Merikle (1993). Theory and measurement in the study of unconscious processes. In M. Davies, & G. W. Humphreys (Eds.), *Consciousness: Psychological and philosophical essays* (pp. 40–57). Oxford: Basil Blackwell.

Rensink, Ronald A. (2000). The dynamic representation of scenes. *Visual Cognition, 7* (1–3), 17–42.

Revonsuo, Antti (1997). How to take consciousness seriously in cognitive neuroscience. *Communication and Cognition, 30* (3/4), 185–206.

Reynolds, John H., & Robert Desimone (2001). Neural mechanisms of attentional selection. In J. Braun, C. Koch, & J. L. Davis (Eds.), *Visual attention and cortical circuits* (pp. 121–136). Cambridge, MA: MIT Press.

Robertson, Lynn C. (1998). Visuospatial attention and parietal function: Their role in object perception. In R. Parasuraman (Ed.), *The attentive brain* (pp. 257–298). Cambridge, MA: MIT Press.

Rosenthal, David M. (1991a). The independence of consciousness and sensory quality. In E. Villanueva (Ed.), *Consciousness* (pp. 15–36). Atascadero, CA: Ridgeview Publishing.

Rosenthal, David M. (Ed.) (1991b). *The nature of mind*. Oxford: Oxford University Press.

Rosenthal, David M. (1991c). Two concepts of consciousness. In D. M. Rosenthal (Ed.), *The nature of mind* (pp. 462–477). Oxford: Oxford University Press.

Rosenthal, David M. (1993a). Higher-order thoughts and the appendage theory of consciousness. *Philosophical Psychology, 6* (2), 155–166.

Rosenthal, David M. (1993b). State consciousness and transitive consciousness. *Consciousness and Cognition, 2*, 355–363.

Rosenthal, David M. (1993c). Thinking that one thinks. In M. Davies, & G. W. Humphreys (Eds.), *Consciousness: Psychological and philosophical essays* (pp. 197–223). Oxford: Basil Blackwell.

Rosenthal, David M. (1995). Multiple drafts and facts of the matter. In T. Metzinger (Ed.), *Conscious experience* (pp. 359–373). Paderborn: Ferdinand Schoeningh.

Rosenthal, David M. (1997). A theory of consciousness. In N. Block, O. Flanagan, & G. Güzuldere (Eds.), *The nature of consciousness: Philosophical debates* (pp. 729–754). Cambridge, MA: MIT Press.

Rosenthal, David M. (2000). Introspection and self-interpretation. *Philosophical Topics, 28* (2), 201–233.

Rosenthal, David M. (2002). Explaining consciousness. In D. Chalmers (Ed.), *Philosophy of mind: Classical and contemporary readings* (pp. 406–421). New York: Oxford University Press.

Rosenthal, David M. (2003). Unity of consciousness and the self. *Proceedings of the Aristotelian Society, 103* (3).

Searle, John R. (1980). Minds, brains and programs. *Behavioral and Brain Sciences, 3* (3), 417–424.

Searle, John R. (1991). Consciousness, unconsciousness and intentionality. In E. Villanueva (Ed.), *Consciousness* (pp. 45–65). Atascadero, CA: Ridgeview Publishing.

Searle, John R. (1992). *The rediscovery of the mind*. Cambridge, MA: MIT Press.

Sellars, Wilfrid (1963). *Science, perception and reality*. Atascadero, CA: Ridgeview Publishing Co.

Shear, Jonathan (Ed.). (1997). *Explaining consciousness: The "hard problem"*. Cambridge, MA: MIT Press.

Shoemaker, Sydney (1996). *The first-person perspective and other essays*. Cambridge: Cambridge University Press.

Singer, Wolf (1996). Neuronal Synchronization: A solution to the binding problem? In R. Llinás, & P. S. Churchland (Eds.), *The mind-brain continuum* (pp. 101–130). Cambridge, MA: MIT Press.

Singer, Wolf (2000). Phenomenal awareness and consciousness from a neurobiological perspective. In T. Metzinger (Ed.), *Neural correlates of consciousness* (pp. 121–137). Cambridge, MA: MIT Press.

Steriade, M., R. Curro Dossi, D. Pare, & G. Oakson (1991). Fast oscillations (20–40 Hz) in thalamocortical systems and their potentiation by mesopontine cholinergic nuclei in the cat. *Proceedings of the National Academy of Sciences of the United States of America, 88* (10), 4396–4400.

Stoerig, Petra (1998). Varieties of vision: From blind responses to conscious recognition. In S. R. Hameroff, A. W. Kaszniak, & A. C. Scott (Eds.), *Toward a science of consciousness II: The second discussions and debates* (pp. 297–308). Cambridge, MA: MIT Press.

Stoerig, Petra, & Erhardt Barth (2001). Low-level phenomenal vision despite unilateral destruction of primary visual cortex. *Consciousness and Cognition, 10*, 574–587.

Stoerig, Petra, & Alan Cowey (1996). Visual perception and phenomenal consciousness. In S. R. Hameroff, A. W. Kaszniak, & A. C. Scott (Eds.), *Toward a science of consciousness: The first Tucson discussions and debates* (pp. 259–278). Cambridge, MA: MIT Press.

Swartz, Robert J. (Ed.). (1965). *Perceiving, sensing, and knowing*. New York: Doubleday and Co.

Tipper, Steven, Louise Howard, & George Houghton (1999). Action-based mechanisms of attention. In G. Humphreys, J. Duncan, & A. Treisman (Eds.), *Attention, space and action: Studies in cognitive neuroscience* (pp. 232–248). Oxford: Oxford University Press.

Tomasello, Michael (1993). On the interpersonal origins of self-concept. In U. Neisser (Ed.), *The perceived self: Ecological and interpersonal sources of self knowledge* (pp. 174–184). Cambridge: Cambridge University Press.

Treisman, Anne (1988). Features and objects. *Quarterly Journal of Experimental Psychology, 40A*, 201–237.

Treisman, Anne (1993). The perception of features and objects. In A. Baddeley, & L. Weiskrantz (Eds.), *Attention: Selection, awareness and control* (pp. 5–35). Oxford: Clarendon Press.

Treisman, Anne (1999). Feature binding, attention and object perception. In G. Humphreys, J. Duncan, & A. Treisman (Eds.), *Attention, space and action: Studies in cognitive neuroscience* (pp. 91–111). Oxford: Oxford University Press.

Treisman, Anne M., & Garry Gelade (1980). A feature-integration theory of attention. *Cognitive Psychology, 12*, 97–136.

Tsotsos, John K., & Sean M. Culhane (2001). From foundational principles to a hierarchical selection circuit for attention. In J. Braun, C. Koch, & J. L. Davis (Eds.), *Visual attention and cortical circuits* (pp. 285–306). Cambridge, MA: MIT Press.

Tye, Michael (1992). Visual qualia and visual content. In T. Crane (Ed.), *The contents of experience: Essays on perception* (pp. 158–176). Cambridge: Cambridge University Press.

Tye, Michael (1995). *Ten problems of consciousness*. Cambridge, MA: MIT Press.

Tye, Michael (1997). The problem of simple minds: Is there anything it is like to be a honey bee? *Philosophical Studies, 88*, 289–317.

Tye, Michael (1998). Inverted earth, swampman, and representationism. In J. E. Tomberlin (Ed.), *Language, mind, and ontology* (pp. 459–477). Malden, MA: Blackwell Publishers.

Tye, Michael (2000). *Consciousness, color, and content*. Cambridge, MA: MIT Press.

Van Gulick, Robert (1992). Time for more alternatives. *Behavioral and Brain Sciences, 15* (2), 228–229.

Villanueva, Enrique (Ed.). (1991). *Consciousness*. Philosophical Issues, 1. Atascadero, CA: Ridgeview Publishing.

Warren, Richard (1992). Global pattern perception and temporal order judgements. *Behavioral and Brain Sciences, 15* (2), 230–231.

Weiskrantz, Lawrence (1997). *Consciousness lost and found: A neuropsychological exploration*. Oxford: Oxford University Press.

Wellman, Henry M. (1988). First steps in the child's theorizing about the mind. In J. W. Astington, P. L. Harris, & D. R. Olson (Eds.), *Developing theories of mind* (pp. 64–92). Cambridge: Cambridge University Press.

Wise, Richard J. S., Sophie K. Scott, S. Catrin Blank, Cath J. Mummery, Kevin Murphy, & Elizabeth A. Warburton (2001). Separate neural sub-systems within "Wernicke's area". *Brain, 124* (1), 83–95.

Young, Andrew W., & Edward H. F. De Haan (1993). Impairments of visual awareness. In M. Davies, & G. W. Humphreys (Eds.), *Consciousness: Psychological and philosophical essays* (pp. 58–73). Oxford: Basil Blackwell.

Zablocka, Teresa, & Boguslaw Zernicki (1996). Discrimination learning of grating orientation in visually deparived cats and the role of the superior colliculi. *Behavioral Neuroscience, 110* (3), 621–625.

Zeki, Semir (1993). *A vision of the brain*. Oxford: Blackwell Scientific Publications.

Index

A

attention 28, 87–88, 141–142, 148–156
  focused and divided 142–144
affordance 114, 117
allocentric space 117–118
Aquinas, St. Thomas 18
Aristotle 17
Armstrong, David 11, 20–21, 25, 44–45, 48, 51, 146
Avicenna 17–18
awareness, *see* thing-awareness *and* fact-awareness

B

Baars, Bernard 81–82
blindsight 5, 8, 12, 57–61, 71, 145
Block, Ned 5–6, 11–12, 28, 123–125

C

camera obscura 135–137
Campbell, John 117
Carruthers, Peter 34–35
Chalmers, David 128–131
chauvinism 28, 123–125
Clark, Austen 3
cognition 7, 23–25, 41–42, 89–91, 137–138
common sense 17–18
conscious sensory states, *see* consciousness, sensory
consciousness
  access 5, 11–13
  creature 10, 26
  degrees of 62–63, 66–71, 87, 95, 127
  intransitive 9–10, 161 Note 54
  introspective 6, 16, 21–22, 25, 31, 45, 50–52, *see also* introspection
  minimal 11
  perceptual 11, 27, 44
  phenomenal 5–6
  sensory ix, 2–4, 13–17, 25–28, 51–52, 62, 66, 69–72, 83–88, 90–92, 95, 111, 119, 126–127, 145–147, 150–156
  state 4–6, 15, 20–22, 25–26, 31–32, 35–41, 45, 47–48, 54–63
  subjective 46
  transitive 9–10, 32–33, 36–37, 40–41, 45, 56–59, 161 Note 54
consciousness theory
  dispositional 34–35
  flat 52–54, 57, 92–94, 98, 135
  higher-order perception, *see* consciousness theory, inner sense
  higher-order thought 25, 31–46, 60, 67, 132–134, *see also* consciousness theory, higher-order
  higher-order 21, 25–27, 32, 35–39, 53–59, 63, 65, 93–94, 98, 135, *see also* consciousness theory, inner sense; consciousness theory, higher-order thought
inner sense 17–22, 25, 31–32, 44–52, 63, 86–87, 110–113, 121, 146–147, *see also* consciousness theory, higher-order

representational ix, 26, 28, 53, 62–63
second sense x, 22–23, 27, 52, 66, 86–98, 124–127, 139, 141–142, 148–156
Wide-Intrinsicality View 40
consumer semantics, *see* teleo-functional semantics

D
Dennett, Daniel 28, 59, 67, 73–78, 93–94, 97, 135–140
Descartes, Rene 7–8
Dretske, Fred x, 25–27, 31, 47, 50, 53–61, 65, 95–98, 135–137

E
egocentric
frames of reference, *see* egocentric, map
map 114, 116–122
space 28, 117–120
eliminativism 100
externalism 54

F
fact-awareness 60–61, 135–136
first-person authority 100–102
functionalism 28, 111, 123–124, 128, 131, 155

G
Gennaro, Rocco 40, 47
Gibson, J. J. 110, 114–118

H
Hard Problem 28, 128–132
higher-order, *see* consciousness theory, higher order; representation, higher-order
homunculus, Cartesian 93–94, 97, 114

I
incorrigibility 102
indexicals 76, 102, 111–112
inner sense, *see* consciousness theory, inner sense
intentional, *see* representation
introspection 18, 45–46, 103–105, 113–114, 121, *see also* consciousness, introspective

J
Jackson, Frank 130–131

K
Kant, Immanuel 19–20
Kinsbourne, Marcel 73–78, 97

L
LaBerge, David 152
language of thought 110–113
liberalism 28, 124–128
Locke, John 8, 18–19
Lycan, William x, 3, 21, 25, 28, 38, 45–46, 48–51, 63, 94, 110–114, 121, 146

M
materialism, *see* physicalism
memory 14–15, 34–35
mental, definition of 3
Millikan, Ruth Garrett 3
Myth of the Given 43

N
Nagel, Thomas 99, 105–110

O
Orwellian editing 78–79, 134, 139–140

P
physicalism 99–101, 105–106, 110, 125–130
point of view 46, 99–104, 106–111, 113, 119–120

Posner, Michael 149
privileged access 99–100, 109

Q
qualia, *see* qualitative character
qualitative character 4–6, 53, 100–101, 105–106, 164 Note 105

R
Rensink, Ronald 143
reportability 133–137, 139–140
representation
  conceptual 69–70, 107, 111–114, 132–138, *see also* cognition
  higher-order 45–46, 48, 65, 87, 91–92, 104, 112–113, 122
  mental 3, 11, 34, 36–37, 47–51, 90, 103–104, 108–110, 121, 126, 136, 157 Note 2, 158 Note 22
  of the present moment 27, 62, 66, 72, 75–86, 126–128, 145, 155
  sensory x, 4, 11, 14, 20, 27, 47, 50, 53, 62, 66–76, 80–82, 85–91, 93–97, 142, 144, 153, *see also* sensation
  virtual 143–144
Rosenthal, David x, 8–10, 15, 25, 28, 32–33, 36–43, 45, 67, 132–135

S
second sense, *see* consciousness theory, second sense
self-conscious sensory states, *see* consciousness, introspective
sensation 13, 23–25, 42–43, 89–91, 137–138, *see also* representation, sensory
sense data 49, 103, 109
sensus communis, *see* common sense
Spot-sight 59–60, 65, 94–97
Stalinesque editing 73–79, 134, 139–140
subjectivity 28, 97, 99–101, 105–106, 109–122, 138–140

T
teleo-functional semantics 3, 34, 136
thing-awareness 60–61, 135
Treisman, Anne 142–143
Twin Earth 111–112
Tye, Michael x, 3, 5–6, 53

U
unconscious sensory states 7–14, 22, 27, 39, 58–62, 72, 82–84, 87–88, 96, 119, *see also* representation, sensory

Z
zombie 126, 128

In the series ADVANCES IN CONSCIOUSNESS RESEARCH (AiCR) the following titles have been published thus far or are scheduled for publication:

1. GLOBUS, Gordon G.: *The Postmodern Brain*. 1995.
2. ELLIS, Ralph D.: *Questioning Consciousness. The interplay of imagery, cognition, and emotion in the human brain*. 1995.
3. JIBU, Mari and Kunio YASUE: *Quantum Brain Dynamics and Consciousness. An introduction*. 1995.
4. HARDCASTLE, Valerie Gray: *Locating Consciousness*. 1995.
5. STUBENBERG, Leopold: *Consciousness and Qualia*. 1998.
6. GENNARO, Rocco J.: *Consciousness and Self-Consciousness. A defense of the higher-order thought theory of consciousness*. 1996.
7. MAC CORMAC, Earl and Maxim I. STAMENOV (eds): *Fractals of Brain, Fractals of Mind. In search of a symmetry bond*. 1996.
8. GROSSENBACHER, Peter G. (ed.): *Finding Consciousness in the Brain. A neurocognitive approach*. 2001.
9. Ó NUALLÁIN, Seán, Paul MC KEVITT and Eoghan MAC AOGÁIN (eds): *Two Sciences of Mind. Readings in cognitive science and consciousness*. 1997.
10. NEWTON, Natika: *Foundations of Understanding*. 1996.
11. PYLKKÖ, Pauli: *The Aconceptual Mind. Heideggerian themes in holistic naturalism*. 1998.
12. STAMENOV, Maxim I. (ed.): *Language Structure, Discourse and the Access to Consciousness*. 1997.
13. VELMANS, Max (ed.): *Investigating Phenomenal Consciousness. Methodologies and Maps*. 2000.
14. SHEETS-JOHNSTONE, Maxine: *The Primacy of Movement*. 1999.
15. CHALLIS, Bradford H. and Boris M. VELICHKOVSKY (eds.): *Stratification in Cognition and Consciousness*. 1999.
16. ELLIS, Ralph D. and Natika NEWTON (eds.): *The Caldron of Consciousness. Motivation, affect and self-organization – An anthology*. 2000.
17. HUTTO, Daniel D.: *The Presence of Mind*. 1999.
18. PALMER, Gary B. and Debra J. OCCHI (eds.): *Languages of Sentiment. Cultural constructions of emotional substrates*. 1999.
19. DAUTENHAHN, Kerstin (ed.): *Human Cognition and Social Agent Technology*. 2000.
20. KUNZENDORF, Robert G. and Benjamin WALLACE (eds.): *Individual Differences in Conscious Experience*. 2000.
21. HUTTO, Daniel D.: *Beyond Physicalism*. 2000.
22. ROSSETTI, Yves and Antti REVONSUO (eds.): *Beyond Dissociation. Interaction between dissociated implicit and explicit processing*. 2000.
23. ZAHAVI, Dan (ed.): *Exploring the Self. Philosophical and psychopathological perspectives on self-experience*. 2000.
24. ROVEE-COLLIER, Carolyn, Harlene HAYNE and Michael COLOMBO: *The Development of Implicit and Explicit Memory*. 2000.
25. BACHMANN, Talis: *Microgenetic Approach to the Conscious Mind*. 2000.
26. Ó NUALLÁIN, Seán (ed.): *Spatial Cognition. Selected papers from Mind III, Annual Conference of the Cognitive Science Society of Ireland, 1998*. 2000.
27. McMILLAN, John and Grant R. GILLETT: *Consciousness and Intentionality*. 2001.

28. ZACHAR, Peter: *Psychological Concepts and Biological Psychiatry. A philosophical analysis.* 2000.
29. VAN LOOCKE, Philip (ed.): *The Physical Nature of Consciousness.* 2001.
30. BROOK, Andrew and Richard C. DeVIDI (eds.): *Self-reference and Self-awareness.* 2001.
31. RAKOVER, Sam S. and Baruch CAHLON: *Face Recognition. Cognitive and computational processes.* 2001.
32. VITIELLO, Giuseppe: *My Double Unveiled. The dissipative quantum model of the brain.* 2001.
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.
34. FETZER, James H.(ed.): *Consciousness Evolving.* 2002.
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.
36. PERRY, Elaine, Heather ASHTON and Allan YOUNG (eds.): *Neurochemistry of Consciousness. Neurotransmitters in mind.* 2002.
37. PYLKKÄNEN, Paavo and Tere VADÉN (eds.): *Dimensions of Conscious Experience.* 2001.
38. SALZARULO, Piero and Gianluca FICCA (eds.): *Awakening and Sleep-Wake Cycle Across Development.* 2002.
39. BARTSCH, Renate: *Consciousness Emerging. The dynamics of perception, imagination, action, memory, thought, and language.* 2002.
40. MANDLER, George: *Consciousness Recovered. Psychological functions and origins of conscious thought.* 2002.
41. ALBERTAZZI, Liliana (ed.): *Unfolding Perceptual Continua.* 2002.
42. STAMENOV, Maxim I. and Vittorio GALLESE (eds.): *Mirror Neurons and the Evolution of Brain and Language.* 2002.
43. DEPRAZ, Natalie, Francisco VARELA and Pierre VERMERSCH.: *On Becoming Aware. A pragmatics of experiencing.* 2003.
44. MOORE, Simon and Mike OAKSFORD (eds.): *Emotional Cognition. From brain to behaviour.* 2002.
45. DOKIC, Jerome and Joelle PROUST: *Simulation and Knowledge of Action.* 2002.
46. MATHEAS, Michael and Phoebe SENGERS (ed.): *Narrative Intelligence.* 2003.
47. COOK, Norman D.: *Tone of Voice and Mind. The connections between intonation, emotion, cognition and consciousness.* 2002.
48. JIMÉNEZ, Luis: *Attention and Implicit Learning.* 2003.
49. OSAKA, Naoyuki (ed.): *Neural Basis of Consciousness.* 2003.
50. GLOBUS, Gordon G.: *Quantum Closures and Disclosures. Thinking-together post-phenomenology and quantum brain dynamics.* 2003.
51. DROEGE, Paula: *Caging the Beast. A theory of sensory consciousness.* 2003.
52. NORTHOFF, Georg: *Philosophy of the Brain. The 'Brain problem'.* n.y.p.
53. HATWELL, Yvette, Arlette STRERI and Edouard GENTAZ (eds.): *Touching for Knowing. Cognitive psychology of haptic manual perception.* n.y.p.

444193

DATE DUE
DATE DE RETOUR

DEC 0 7 2004
JAN 1 2 2005
JAN 1 3 2005

CARR McLEAN
38-296

TRENT UNIVERSITY
0 1164 0496653 7

A major obstacle for materialist theories of the mind is the problem of sensory consciousness. How could a physical brain produce conscious sensory states that exhibit the rich and luxurious qualities of red velvet, a Mozart concerto or fresh-brewed coffee?

Caging the Beast: A Theory of Sensory Consciousness offers to explain what these conscious sensory states have in common, by virtue of being conscious as opposed to unconscious states. After arguing against accounts of consciousness in terms of higher-order representation of mental states, the theory claims that sensory consciousness is a special way we have of representing the world. The book also introduces a way of thinking about subjectivity as separate and more fundamental than consciousness, and considers how this foundational notion can be developed into more elaborate varieties. An appendix reviews the connection between consciousness and attention with an eye toward providing a neuropsychological instantiation of the proposed theory.

John Benjamins Publishing Company

ISBN 90 272 5182 7 (Eur) / 1 58811 391 4 (US)