A HOLISTIC APPROACH TO REPRESENTATIONALISM

by

JACOB BERGER

A dissertation submitted to the Graduate Faculty in Philosophy
in partial fulfillment of the requirements for the degree of
Doctor of Philosophy, The City University of New York

2013

This manuscript has been read and accepted for the Graduate Faculty in Philosophy in satisfaction of the dissertation requirement for the degree of Doctor of Philosophy.

Michael Devitt

Date — Chair of Examining Committee

Iakovos Vasiliou

Date — Executive Officer

Jesse Prinz, Advisor

Stephen Neale

Tony Ro

Richard Brown

Supervisory Committee

THE CITY UNIVERSITY OF NEW YORK

Abstract

A HOLISTIC APPROACH TO REPRESENTATIONALISM

by

Jacob Berger

Advisor: Professor Jesse J. Prinz

Perhaps the most promising account of the qualitative character of experience available is representationalism—the view that the qualitative character of a mental state is identical with (or supervenes on) that state's representational properties. According to representationalism, for example, the reddish qualitative character of a perception is (or is determined by) the property of the state's qualitatively representing red. But representationalism is incomplete without an account of how experiences represent what they do—that is, an account of the psychosemantics of qualitative content. To date, most representationalists have endorsed versions of so-called tracking theories of content, according to which a state represents a property just in case the state tracks that property. Such views are atomistic insofar as a state's content does not depend on its relations to other mental states. Versions of representationalism which depend upon such atomistic psychosemantics are, however, open to criticism. Some representationalists have therefore concluded that qualitative representation is primitive or resists reductive explanation. But this reaction may be too hasty. This dissertation develops a form of reductive representationalism according to which qualitative content is individuated in a holistic way.

To develop this view, Chapter 1 addresses introductory issues regarding qualitative character and representation. Chapter 2 argues that standard forms of representationalism of the sort defended by Fred Dretske, William Lycan, and Michael Tye fail primarily because of their

atomistic approach to qualitative content. Recently, some representationalists have offered more sophisticated versions of the view, principally to accommodate phenomena such as undetectable quality inversion. Chapter 3 argues that these more complex accounts—including Sydney Shoemaker's dispositionalist representationalism and David Chalmers's Fregean representationalism—are unworkable. One might think that the failures of these views suggest that the chief rival to representationalism, the traditional qualia theory which holds that there are nonrepresentational qualitative aspects of perceptions, is correct. Chapter 4 argues that such a qualia theory is problematic because it cannot provide an account of our knowledge of qualia, even from the first-person perspective.

In light of these considerations, Chapter 5 proposes a version of representationalism wherein qualitative content is individuated in a holistic way. This holistic theory of qualitative content—what is dubbed here 'perceptual-role semantics'—builds upon a burgeoning theory of qualitative character, versions of which have been defended by, among others, Austen Clark, David Lewis, David Rosenthal, and Shoemaker. On the view developed, a qualitative state's content is determined by its relative location in a space of states that matches the corresponding quality spaces of perceptible properties to which those qualitative states provide access. For example, an experience of red represents red because the experience occupies a location within a space of experiences of color that corresponds to the location occupied by red within the quality space of colors. The resultant holistic version of representationalism avoids the problems that plague other versions of it, resolves a host of philosophical puzzles about qualitative character, fits with a range of recent empirical findings about perception, and opens the phenomena up to fruitful further study.

Acknowledgements:

Aspects of this dissertation were presented at the Association for the Scientific Study of Consciousness, the Second and Third Annual Online Consciousness Conferences, the CUNY Cognitive Science Symposium, Cambridge University, New York University, Byrn Mawr College, and City College, CUNY. I thank these audiences for their discussions of this material.

This dissertation has benefited from the kind input of many people. A number of fellow students, colleagues, and professors have discussed and commented on parts of this work, including Mark Alfano, Valtteri Arstila, Ned Block, Richard Brown, Jennifer Corns, Ryan DeChant, Michael Devitt, Willem de Vries, Kyle Ferguson, Daniel Harris, Benj Hellie, Hasen Khudairi, Alex Kiefer, Thomas Kivatinos, Michal Klincewicz, Uriah Kriegel, Jed Lewinsohn, Pete Mandik, Stephen Neale, Adam Pautz, Ben Phillips, Daniel Shargel, Joshua Shepherd, Jake Quilty-Dunn, and Josh Weisberg. I am indebted to Tony Ro and the members of his cognitive neuroscience group—especially Tatiana Emmanouil and Marjan Persuh—for welcoming me into their lab. I owe particular thanks to Myrto Mylopoulos and David Pereploytchik for their invaluable support throughout this process. I wish also to acknowledge my deep debt to David Rosenthal, whose personal advice and contributions to the field have shaped and influenced this dissertation in myriad ways.

Special thanks are due to my dissertation advisor, Jesse Prinz, who has been an intellectual inspiration and unending source of guidance and encouragement.

I am grateful to all of my patient friends and family—and especially to my parents, Drs. Elizabeth and Marc Berger—for their love and good humor.

I dedicate this dissertation to my wife, Kristen, who puts it all in perspective.

A HOLISTIC APPROACH TO REPRESENTATIONALISM

Table of Contents

CHAPTER 1: REPRESENTATION AND QUALITATIVE CHARACTER

§1.1 Introduction

Perhaps nothing seems more central to our mental lives than the qualities of our experience—the fact that at least some of our mental states exhibit qualitative character. Examples of qualitative character include the reddish quality of seeing a tomato, the lemony quality of tasting lemonade, and the painful quality of experiencing a bee sting. To adapt Thomas Nagel's (1974, p. 436) famous expression, these are the properties in virtue of which there is something that it's like to be in a conscious perception or sensation. As Daniel Dennett notes, there are various expressions for these qualitative dimensions of mental states such as "'raw feels', 'sensa', 'phenomenal qualities', 'intrinsic properties of conscious experiences', the 'qualitative content of mental states' and of course 'qualia'" (1991, p. 372). I will call these properties mental qualities.[1] In this terminology, a visual sensation of a blueberry exhibits a mental quality that I'll call mental blue.

Common sense holds that mental qualities are manifest in consciousness. Nonetheless, there is much disagreement among philosophers about what role, if any, mental qualities play in psychological functioning. For various reasons that I'll discuss later in this chapter, many theorists believe that mental qualities resist any informative explanation in functional or naturalistic terms (e.g., Nagel 1974, Kripke 1980, Jackson 1982, Block 1990, Chalmers 1996, Levine 2001). Mental qualities thus lie at the heart of many traditional philosophical puzzles and the study of mental qualities remains one of the core issues in the philosophy of mind.

[1] I use the expression 'mental qualities' and not the far more common term 'qualia' because some theorists have repurposed the latter to refer to the qualities of objects (such as the redness of an tomato), not mental properties (e.g., Tye 2000, p. 49). I borrow the expression 'mental qualities' from David Rosenthal (2005).

Despite the lack of consensus about the nature of mental qualities, many historical and contemporary thinkers distinguish mental qualities from the representational properties of the mind (e.g., Reid 1764/1997, Kant 1787/1998, Block and Fodor 1972, Nagel 1974, Peacocke 1983). Cognitive states, such as such as beliefs, desires, hopes, and doubts, are often thought to be paradigmatic examples of mental representations. For example, most would agree that Jack's belief that the dog barks represents something like the state of affairs that the dog barks.[2] In short, a mental representation is a psychological state that stands for something so that the state can operate in psychological functioning pertaining to that thing. Following Franz Brentano (e.g., 1874/1973, p. 88), states that represent or are about things are said to exhibit intentionality.[3] The idea that at least some mental states are representational or intentional is, for the most part, not in question; mental representations are stock-in-trade to both contemporary philosophical theorizing and cognitive science (see, e.g., Stich and Warfield 1994, Pitt 2008).

But many philosophers assume that mental qualities are not representational. For example, Thomas Reid claims that while perceptions are representational—they have corresponding objects—sensations are not. Reid writes,

> I shall add an observation concerning the word feeling. This word has two meanings. First, it signifies the perceptions we have of external objects, by the sense of touch. When we speak of feeling a body to be hard or soft, rough or smooth, hot or cold; to feel these things, is to be perceive them by touch. They are external things, and that act of the mind by which we feel them, is easily distinguished from the objects felt. Secondly, the word feeling is used to signify the same thing as sensation, which we have just now explained; and, in this sense, it has no object; the feeling and the thing felt are one and the same (1785/2002, p. 38, emphasis his).

More recently, Brian Loar notes,

[2] For now, I will remain neutral about what it is exactly that states such as beliefs represent. I will have more to say about this in chapter 5.

[3] Brentano borrowed the term 'intentionality' from Scholastic philosophers. The term was drawn from the medieval Latin 'intentio' and was used to describe ideas or mental representations (Crane 2005, p. 438).

> According to the standard view qualia are not in themselves, not intrinsically, representational or intentional…. Like paint on canvas, qualia are individuated, on this view, independently of their representational or referential properties, and—again like paint—individuated independently even of their purporting to represent, independently of their having representational properties even in Brentano's sense. We can say that what the standard view defends are raw qualia (2003, p. 78, emphasis his).

Call the view that mental qualities are mere "raw" feels that do not represent anything the traditional qualia view. According to a traditional taxonomy of mental states, some states such as beliefs are representational but do not exhibit mental qualities, some states such as pains exhibit mental qualities but are not representational, and some states such as emotions are both representational and exhibit mental qualities. I will call states that exhibit mental qualities only sensations. Though 'sensation' is often reserved for bodily sensations such as pains, tickles, and itches, we can equally speak of perceptual sensations, which are sensations of external properties such as the blueness of a blueberry (cf. Clark 1993, pp. 3-4). A visual sensation of blue, then, exhibits mental blue only. Perceptions, by contrast, are usually thought to involve some representational properties. On this classification, a perception that there is blue exhibits both mental blue and represents something like the state of affairs that there is blue.

Recently, the traditional qualia view has been challenged. There is increasing evidence, for example, that qualitative states such as perceptions and sensations are representational. Indeed, perhaps the most promising view in the contemporary philosophy of perception is the view known as representationalism, which holds that the qualitative character of a state is identical with (or supervenes on) that state's representational properties (e.g., Armstrong 1968, Harman 1990, Dretske 1995, Tye 1995, Lycan 1996, Byrne 2001, Thau 2002, Jackson 2007,

Crane 2009, Hill 2009, Speaks 2009, Pautz 2010).[4] As Fred Dretske puts it, according to representationalism, "[a]ll mental facts are representational facts" (1995, p. xiii).

Representationalism comes in many varieties (for a survey, see, e.g., Chalmers 2004). The view that mental qualities are identical with (or supervene on) representational properties is often referred to as strong representationalism. According to what is often known as weak representationalism, by contrast, qualitative states exhibit representational properties. Weak representationalism is, of course, compatible with qualitative states exhibiting nonrepresentational qualia too. And, as David Chalmers observes, "Weak representationalism (at least about visual experiences) is extremely plausible, and rarely denied" (2004, p. 344).[5] Uses of 'representationalism' throughout will refer to strong representationalism, unless otherwise noted.[6]

Most representationalists claim that the qualitative character of a state is identical with (or supervenes on) that the property of a state's representing (in the right qualitative way) what I will call perceptible properties.[7] Common sense holds that objects exhibit a range of perceptible

[4] Representationalism often goes by various names, such as intentionalism or representationism. The version of representationalism according to which qualitative character supervenes on representational properties is a comparatively weaker than the view that they are identical because, as Tye observes, "this supervenience thesis leaves open the further question as to the essential nature of qualia" (2007, §7). In this dissertation, I will argue for a version of the identity thesis, though I am happy to settle for a version of the supervenience view. For the most part, I will remain neutral between these versions of the view.

[5] References to Chalmers's papers reprinted in Chalmers (2010) will reference the year of the original publication, but the page(s) of the paper reprinted in his (2010).

[6] I will also draw distinctions among other relevant varieties of the view as this dissertation progresses.

[7] I say that mental qualities are (or are determined by) the way the states represent perceptible properties in the right qualitative way because, naturally, one can represent those properties in nonqualitative ways. For example, one can believe that there is perceptible red present. Thus representationalists have offered various explanations of the difference between qualitative and cognitive representation (for a survey of such views, see, e.g., Bourget and Mendelovici

properties such as colors, smells, and tastes. A red tomato, for example, exhibits the perceptible color red. As I'll discuss shortly, many theorists have denied that perceptible properties exist. But it is clear that if mental qualities and perceptible properties both exist, they are distinct kinds of properties. For one thing, perceptible properties are properties of mind-independent objects, whereas mental qualities are properties of mental states.[8] For clarity, I will refer to the properties of objects as perceptible properties and to the corresponding properties of mental states as mental qualities. In my terminology, a ripe tomato exhibits perceptible red and a visual perception of that tomato exhibits mental red. According to representationalism, mental red simply is (or is determined by) the property of the state's representing (in the right way) perceptible red.

There are several reasons to think that qualitative states are representational. First, it seems clear that qualitative states can be assessed for accuracy (see, e.g., Pautz 2010a, p. 269; Siegel 2010, p. 31; Nanay forthcoming, p. xxx). If I am in a state that exhibits mental red, then it would seem that my state is accurate just in case perceptible red is present and inaccurate otherwise. Similarly, whenever one's qualitative states change, it appears that the way things seem to one change as well (see, e.g., Byrne 2001, p. 211; Thau 2002, pp. 32-33; Schellenberg 2011, pp. 719-720; for discussion, see, e.g., Lycan 2006, §3.4 or Crane 2009b, pp. 484-486). That is, if I shift from being in a state that exhibits mental red to a state that exhibits mental green, then at first things seem to me to be perceptible red and then things seem to me to be perceptible green. Third, qualitative states connect with our perceptual thoughts—that is, our thoughts about what we are currently perceiving or sensing—in a particular way (see, e.g.,

---

forthcoming). I will have more to say about what distinguishes qualitative states and cognitive states in chapter 5.

[8] Many historical and contemporary theorists thus maintain that the predicates for perceptible properties such as 'red' ambiguously refer both to perceptible properties and to their corresponding mental qualities (e.g., Reid 1785/1969, p. 244; Peacocke 1983, pp. 20-21; Clark 2000, p. 2; Levine 2001, p. 179, n. 5).

Sellars 1967, p. 17; Pautz 2010a, p. 266; Nanay forthcoming, p. xxx). If I am in a state that exhibits mental red, then I am disposed to form the perceptual thought that perceptible red is present. All of these facts, which I will discuss in detail, suggest that some form of representationalism is true.

Nevertheless, representationalism remains contentious. Proponents of the traditional qualia view—so-called qualophiles—typically deny representationalism by offering putative evidence of nonrepresentational qualia (e.g., Peacocke 1983, Block 2003, Burge 2003). And it does not appear that the controversy will soon abate. As Ned Block observes,

> The greatest chasm in the philosophy of mind—maybe even all of philosophy—divides... on whether there is anything in the phenomenal character of conscious experience that goes beyond the intentional, the cognitive, the functional (2003, p. 533).[9]

This dissertation is a defense of a version of strong representationalism. And despite qualophiles' resistance to this view, there are many reasons to welcome some form of it. First, representationalism offers a unified account of the mind, insofar as it upholds Brentano's (1874/1973, p. 88) thesis that the mark of the mental is intentionality (e.g., Crane 2009, pp. 474-475). Additionally, since many are confident that we can provide a naturalistic account of mental representation, representationalism has the potential to provide a fully naturalistic account of qualitative character (e.g., Cutter and Tye 2011, p. 90).

In order to deliver on the promise of naturalizing qualitative character, representationalists must, of course, provide a naturalistically credible account of how qualitative representation works. It is incumbent upon representationalists to explain why, for example, mental red represents perceptible red, as opposed to another color, or some other perceptible property altogether. Thus, in this dissertation, I'll propose a novel account of the

[9] References to Block's papers reprinted in Block (2007a) will reference the year of the original publication, but the page(s) of the paper reprinted in his (2007a).

representational character of mental qualities. This theory will form the basis of a new form of representationalism, an account that I'll argue explains mental qualities' role in our psychological lives, resolves a host of philosophical puzzles about them, fits with a range of recent empirical findings about perception, and opens them up to fruitful empirical study.

Before moving on, I should make one important caveat about how these issues relate to consciousness. Representationalists explain qualitative character of states such as perceptions in terms of representation. But, as many representationalists have observed, perception can and often does occur in the absence of consciousness, as in cases of masked priming or blindsight (e.g., Tye 2000, pp. 62-63).[10] Consequently, it is important to distinguish two questions about the nature of mental qualities. Since states with different mental qualities can all occur consciously, there is what, following Adam Pautz (2010, p. 333), we can call the general question—namely, the question of what explains why a given qualitative state is conscious. But we must distinguish the general question from what Pautz calls the quality question: What explains why a conscious qualitative state exhibits the particular qualitative character it does?

To give a complete theory of conscious qualitative character, whatever theory of qualitative representation representationalists ultimately offer must be paired with an explanation of what makes the relevant qualitative states conscious. And there are a variety of theories of

[10] A wealth of experimental results provides evidence for nonconscious perception. Though much perception is conscious, there is substantial evidence that we can also perceptually discriminate stimuli without being aware that we perceive those stimuli (e.g., Breitmeyer et al 2004, Ro et al 2009). For example, in a typical metacontrast-masking study, participants are briefly presented with a stimulus, which is immediately followed by a non-overlapping mask that renders the stimulus invisible to consciousness (see, e.g., Breitmeyer & Öğmen 2000). Though participants report that they did not see the stimuli, they nevertheless can be primed by these stimuli in various ways that can be behaviorally detected. The most reasonable explanation of this phenomenon is that participants subliminally perceive these masked stimuli. Since these subliminal perceptions play the same kinds of functional roles as their conscious counterparts, it is arguable that these are cases of nonconscious mental states.

consciousness on the market including attentional theories (e.g., Prinz 2012), higher-order theories (e.g., Lycan 1996, Rosenthal 2005), global-workspace theories (e.g., Baars 1988), and so on. Tye (2000, p. 62), for example, holds that a qualitative state is conscious just in case it is suitably poised for impact on cognition. I cannot settle the issue of consciousness here. But whatever theory explains how qualitative representations are conscious, my goal in this dissertation is to develop and defend a theory of how qualitative states represent what they do—that is, to provide an answer to the quality question.

The plan for this dissertation is as follows: In the remainder of this chapter, I'll place my version of representationalism in context by discussing the nature of mental representation generally. I'll subsequently discuss many of the reasons why most theorists have denied that mental qualities are representational. Chapters 2, 3, and 4 challenge several alternative accounts of mental qualities, proposals that undermine the view that I'll develop in the fifth and final chapter. Chapter 5 then develops my particular version of representationalism, defends it from various objections, and shows how it clarifies some debates in the contemporary philosophy of mind.

§1.2 Mental Representation in General

Mental representations figure prominently in much contemporary philosophical theorizing and cognitive science. For example, many theorists have argued that we can account for the explanatory success of ordinary folk psychology—which explains peoples' behavior in terms of causally efficacious mental states such as beliefs and desires—by understanding these folk-

psychological states to be mental representations (e.g., Fodor 1987, Dretske 1988).[11] This has come to be known as the Representational Theory of Mind (see, e.g., Sterelny 1990).

So what is it for a mental quality—or any mental property for that matter—to be representational? I confess that I do not have a complete account of representation in general, nor even an account of mental representation in particular. I suspect that no one can put forward such an account at present. But, as a first attempt at a characterization, I propose that a mental representation is a mental state that is about or stands for something so that the state can operate in psychological functioning pertaining to that thing. Recall Jack's belief, which represents something like the state of affairs that the dog barks. Because the belief represents that state of affairs in Jack's mind, the belief can play a variety of psychological functions and can issue in various behaviors related to that state of affairs. Jack can infer from this belief that the dog is hungry and this belief can cause him to feed the dog.

In contemporary treatments of mental representation, mental representations are often said to exhibit representational content. Some refer to representational content as intentional or representative content, and I will sometimes simply use 'content' for short. For example, most hold that Jack's belief that the dog barks exhibits the content *that the dog barks*.[12] In order to

[11] In addition to explaining the efficacy of folk psychology, representations have been posited by cognitive scientists to explain many other phenomena such as, for example, our capacities to learn and understand natural language (see, e.g., Chomsky 1965). But it is arguable that many of the representations posited to explain these capacities are, to borrow Daniel Dennett's (1969, p. 93) expression, subpersonal—that is, nonmental states of the organism. Since in chapter 5 I will be defending a view regarding the nature of mental qualities, which (if they exist) are personal-level mental properties, I will restrict my attention to what it is that makes a personal-level mental state a representation.

[12] I will denote the representational contents of cognitive states such as beliefs by enclosing them within this '*' notation.

understand the proposal that mental qualities are representations, it will be helpful to further elucidate exactly what content amounts to.[13]

Many theorists say that a state's content is what the state represents. For example, Uriah Kriegel writes, "A mental state's representational content is what it represents or purports to represent" (2006, p. 59). But phrases such as 'what a state represents' are ambiguous. On one hand, 'what a state represents' can refer to the things that representations are about or represent there to be. Some theorists clearly have this understanding of 'content' in mind: For example, Jesse Prinz writes that "those things to which [states] refer, I call their representational contents" (2002, p. 4). On this usage of 'content', the content of the belief that the dog barks is something like the state of affairs that the dog barks. On the other hand, since a representation exhibits or instantiates a content, 'what a state represents' can instead refer to the way in which a state represents what it is that it is about or concerns. Other theorists use 'content' in this way. Consider, for instance, Colin McGinn's rather standard way of introducing the notion of content: "The way in which experience represents the world constitutes its content" (1989, p. 58).[14]

But even phrases such as 'the way a state represents things' are ambiguous. In order to expose this ambiguity, consider the fact that cognitive states such as beliefs, desires, hopes, and doubts involve at least two distinct psychological components. Common sense as well as most theorists maintain that such states involve holding what I'll call a mental attitude toward a

[13] The following discussion about how to understand content is much influenced by Walter Hopp's (2011, pp. 8-12) treatment of the issue.

[14] As Hopp observes (2011, p. 11), some theorists are explicitly inclusive in their ways of using 'content'. For instance, Hopp quotes Christopher Peacocke, who writes, "Henceforth, I use the phrase 'content of experience' to cover not only which objects, properties and relations are perceived, but also the ways in which they are perceived" (2001, p. 241).

representational content (e.g., Vendler 1972, chapter 3; Searle 1983, p. 6).[15] For example, Jack's belief that the dog barks involves holding the belief attitude toward the content *that the dog barks*.[16] Naturally, one can hold differing attitudes towards this content and can hold the same attitude towards differing contents. One can not only believe that the dog barks, but also desire that the dog barks, wonder whether the dog barks, hope that the dog barks, and so on. And one can believe that the dog barks, believe that it's raining, believe that the tomato is red, and so forth. For states that exhibit mental attitudes, then, a state's mental attitude constitutes an aspect of the way in which the state represents what it is about. The belief that the dog barks and the hope that the dog barks are two ways of representing the state of affairs that the dog barks.

Once we distinguish attitudes from contents, how are we to understand contents? Though mental attitudes are aspects of representational states, they are not the properties in virtue of which representational states represent things. One cannot, of course, simply believe or hope full stop. But if there were just a belief (full stop), such a state would not yet be about anything. In order for a state that exhibits a mental attitude to represent something, one must hold a mental attitude toward something—and what one holds mental attitudes towards are contents. Contents are thus the properties of states in virtue of which states represent things.

---

[15] Bertrand Russell called states that exhibit this content/attitude structure propositional attitudes (e.g., 1948, pp. 167-168). But this name is misleading for several reasons. First, Russell used this terminology because he conceived of representational contents as propositions, abstract objects that determine the truth-conditions of the mental states that exhibit those contents (many still hold this view of contents; see, e.g., Byrne 2005, p. 232). But one might reject this way of characterizing content for nominalistic reasons. If one wishes to avoid abstract objects (as I do), one can instead understand contents simply to be mental properties that type representational states (for this kind of view, see, e.g., Burge 2010, pp. 34-42). Moreover, the expression 'propositional attitude' elides the distinction between a state's attitude and its content. This distinction, however, is clearly embedded in folk psychology and central to theorizing about the mind.

[16] The representational contents of cognitive states are typically understood to be the properties denoted by the that-clause portions of descriptions of them and their attitudes are denoted by the psychological verbs preceding the that-clauses.

To further illustrate the difference between the above two uses of 'content', consider the fact that two distinct beliefs can be about the same thing, but represent that thing in different ways. For example, my belief that the dog barks and my belief that my pet barks both exhibit the belief attitude and concern the same state of affairs—namely, the state of affairs that my pet dog, Fido, barks. But these beliefs are nonetheless distinct because they represent that state of affairs in different ways. One belief represents Fido as a dog and the other belief represents him as a pet. The beliefs represent the same content (understood as what the states are about) via distinct contents (understood as the properties in virtue of which the states represent what they are about).[17]

I will reserve 'content' to refer to the properties of mental states in virtue of which they represent what they do. Instead of 'content', I'll call what a state is about its target.[18] On this usage, states do not represent their contents—rather, they represent their targets by exhibiting contents. For example, the target of the belief that the dog barks is Fido, though the way that I represent Fido is via a state that exhibits the belief attitude towards the content *that the dog barks*.

This distinction between targets and contents is crucial.[19] Once we draw this distinction, we can recognize that there is an open question whether all mental representations have the same kind of psychological structure. Holding a mental attitude towards a content is one way to

[17] The distinction is most vivid if one conceives of contents as abstract objects (see fn. 15). The belief that the dog barks is not about an abstract object, but rather Fido (a concrete dog), even if my belief represents Fido by exhibiting an abstract content.
[18] I borrow this term from Robert Cummins (1996, p. 6), although he uses it for somewhat different purposes.
[19] Brentano (1874/1973, p. 88) himself distinguished contents from targets, but called targets intentional objects. I prefer 'target' because, although 'object' can be understood in a very broad way, targets need not be objects. I will argue in chapter 5 that mental qualities represent properties, not objects.

represent things. But I'll argue in chapter 5 that some representational states do not exhibit this content/attitude structure. Though cognitive states exhibit this structure, I'll argue that sensations, which exhibit mental qualities only, do not exhibit anything like mental attitude.

The idea that states such as sensations and perceptions are representational is not new. Fifty years ago, Wilfrid Sellars (1963) urged that, though sensations are representational, their representational character should not be confused with the representational nature of cognitive states such as beliefs. Sellars writes that "sensations have what I call a pseudo-intentionality, which is easily mistaken for the intentionality of the cognitive order" (1963, p. 46). And Sellars warns that many philosophical missteps ensue when one fails to distinguish these two sorts of representation. This dissertation can be seen as a development of that idea—an explanation of the unique way in which mental qualities represent what they do and a defense of such a view from various challenges.

But Sellars denied that sensations are genuinely intentional because he, like most at his time, assumed that the only representational states are cognitive states such as beliefs. That is, Sellars assumed that all intentional states are conceptual.[20] Tradition holds that cognitive states are conceptual insofar as their contents are structured by term-sized representations—concepts—which suitably combine with other concepts to form those states' contents (e.g., Locke 1700/1975, Hume 1739/2000, Fodor 1975). For example, the belief that the dog barks involves the concepts DOG and BARK,[21] which, properly organized, form the belief's content. I'll reserve 'cognitive state' for states that exhibit this kind of psychological structure—that is, a mental attitude toward a conceptual content.

[20] For example, Sellars writes that "the intentional is that which belongs to the conceptual order" (1967, p. 23).
[21] I follow Jerry Fodor's convention of capitalizing terms when they refer to concepts (see, e.g., Fodor 1998a, p. xii). For example, 'DOG' refers to the concept of a dog.

Nowadays, however, it is common to distinguish several varieties of representational states. Though there is considerable debate about what exactly it is for a state to be conceptual or to exhibit conceptual content (see, e.g., Heck 2000, Byrne 2005, Speaks 2005), many theorists now maintain that some mental representations exhibit so-called nonconceptual content (see, e.g., Evans 1982, Peacocke 1992, Dretske 1995, Tye 1995, Fodor 2007, Burge 2010; for a survey, see, e.g., the essays collected in Gunther 2003). Since concepts and the cognitive states in which they figure are typically considered to be psychologically sophisticated, one of the main reasons to introduce nonconceptual content is, as Tim Crane puts it, "to identify a form of mental representation which is in some ways more primitive, more basic than belief" (2009a, p. 466). On this kind of view, while all conceptual states are intentional, not all intentional states involve conceptual content.

There is a variety of accounts of nonconceptual content on offer and, as Paul Coates (2007, p. 12) observes, nonconceptual contents have been posited to explain many diverse phenomena. For example, Martin Davies (1987, p. 445) has argued that the tacit knowledge of syntax implicated in language processing involves nonconceptual content. These are not the sorts of representations that interest us here, however, because they do not figure in sensation or perception. Nevertheless, some representationalists have claimed that what is distinctive of qualitative states such as perceptions and sensations is that they exhibit nonconceptual content (e.g., Dretske 1995, Tye 1995, Crane 2009a). And, in chapter 5, I will develop a theory of qualitative character according to which mental qualities are (or are determined by) a certain kind of nonconceptual content that I'll call qualitative contents.[22]

[22] I remain neutral whether or not other sorts of nonqualitative nonconceptual content exist.

If representationalists are correct that mental qualities are (or are determined by) a kind of content, then it is incumbent upon them to offer a theory of why a particular mental quality represents what it does. In other words, representationalists must provide an account of the psychosemantics of qualitative contents—an account of the contents of qualitative states which are identical with (or which determine) mental qualities. As most will recall, there was considerable debate in the 1980's about how to naturalize intentionality. At that time, the focus was not as much on the content of qualitative states, but on the content of cognitive states such as beliefs and desires. It is quite plausible, however, that the contents of qualitative states are of a different kind than the contents of conceptual states—so it is open that whatever theory of the semantics of the conceptual contents of cognitive states we give will be different than the theory of the nonconceptual contents of qualitative states (cf. Cutter and Tye 2011, p. 106, fn. 3).

Out of the debates of the 80's, there emerged two broad approaches to conceptual content. On one hand, many philosophers developed atomistic accounts of conceptual content, such as causal-covariational theories, teleosemantics, and asymmetric-dependence theory (e.g., Stampe 1977, Dretske 1981, Papineau 1984, Fodor 1987). According to atomistic theories, a thought's conceptual content is determined by the thought's standing in an appropriate connection to a feature of the world. For example, the thought that there is a dog represents the dog because the thought is connected to the dog in the right way. Such views are atomistic insofar as a state's content does not depend on its relations to other mental states. Indeed, it is consistent with atomism that a creature be capable of exhibiting one thought only.

By contrast, other philosophers endorsed holistic accounts of conceptual content, such as inferential- or conceptual-role semantics ("CRS") (e.g., Sellars 1963, Block 1986, Harman 1987), according to which a thought's content is determined by its inferential relations to other

thoughts (and, on some versions, its causal relations to behavior or elements in the environment as well). These views are holistic insofar as a state's content depends on its relations to other mental states. In order for a creature to have any thoughts at all, the creature must be capable of exhibiting a repertoire of thoughts that stand in appropriate inferential relations to one another.

To date, most representationalists simply import an atomistic theory of content to explain the contents of qualitative states (e.g., Dretske 1995, Tye 1995, Lycan 1996, Hill 2009; for a survey of such views, see, e.g., Bourget and Mendelovici forthcoming)—a holistic account of qualitative content has not been adequately explored. Such atomistic versions of representationalism appear to reduce representation in a naturalistic fashion by specifying some relation (which can be characterized in nonrepresentational physical terms) between a qualitative state and a perceptible property in virtue of which that state represents that property—and thereby offer a naturalistic account of qualitative character.

Recently, however, forms of representationalism that depend on atomistic theories of content have been criticized on various grounds (e.g., Chalmers 2006, Pautz 2006, Bourget 2010). Instead of rejecting representationalism, some have concluded that the contents of qualitative states are primitive, nonreductive, or otherwise resistant to informative explanation (e.g., Bourget 2010, Pautz 2010). But this reaction is too hasty.

Most representationalists simply assume an atomistic semantics for the view, I think, largely because holistic accounts of conceptual content such as CRS are often thought to face insuperable difficulties (see, e.g., Fodor and Lepore 1992). While I believe that holistic theories of conceptual content have much to recommend for them, I will remain neutral in this dissertation about what theory explains the contents of conceptual states. But fundamental to any account of qualitative states is an explanation of our perceptual discriminatory behaviors—

and in this dissertation I'll argue that the theory of the contents of qualitative states that best explains these and other phenomena is holistic.

In particular, I'll propose a holistic theory of the contents of qualitative states that I'll call perceptual-role semantics ("PRS"). This theory, I'll argue, provides the foundation for promising reductive version of representationalism—a view that I'll call holistic representationalism. In chapter 5, I'll argue that, however conceptual representation works,[23] what I'll call conceptual representation—or representation via cognitive states—is a very different kind of representation than what I'll call qualitative representation, which is representation via nonconceptual mental qualities. I'll argue there that mental red represents perceptible red in a quasi-holistic way, which is quite unlike the way in which the belief that there is perceptible red or the concept RED represents it.

However it is that states represent their targets, it is a well-known feature of representations that their targets need not exist. For example, a target of the belief that the unicorn has a horn is a unicorn, even though there are no unicorns. Brentano (1874/1973, p. 88) famously described this aspect of representations as the intentional inexistence of the objects of representations. Whether or not mental qualities are a species of representation, this is an interesting feature of representations generally that any theory must respect and explain.

It is thus central to the notion of representation that representations can misrepresent (e.g., Stich and Warfield 1994, pp. 6-7). For example, I may believe that it is raining and thereby represent the state of affairs that it is raining, even if it is not in fact raining. But if representations can go wrong, then there must be some standard by which we judge their

[23] It is, of course, beyond the scope of this dissertation to defend a particular view about the nature of conceptual content. For discussion of the pros and cons of various theories of conceptual content, see, e.g., Stich and Warfield 1994.

correctness. As David Pitt observes, it is in the nature of a representation to be capable of being "evaluated with respect to properties like consistency, truth, appropriateness and accuracy" (2008, §1). In other words, representations exhibit conditions under which they are veridical and conditions under which they are falsidical. Following Tyler Burge (2010, p. 9), I call these conditions veridicality conditions.[24] A representation's veridicality conditions are met just in case the target that the state represents obtains.

Perhaps the best-understood kinds of veridicality conditions are truth conditions. Consider again Jack's belief that the dog barks, which exhibits the representational content *that the dog barks*. If the belief's target state-of-affairs obtains, then the belief's content is true and false otherwise. These are the state's truth conditions—the conditions under which the belief's content is true or false. The belief that the dog barks represents the state of affairs that the dog barks and its content is thus true when it is barking and false when it is not barking.

We can, of course, not only say that a belief's content is true or false, but also that the belief itself is true or false. But there are other kinds of cognitive states—such as desires or intentions—that we do not describe as being either true or false. We might describe a desire as fulfilled or unfulfilled, but we do not describe it as true or false. That is to say, though states with some mental attitudes can be described as being the bearer of truth and falsity, some states with other attitudes cannot. Whatever a state's attitude, however, cognitive states nonetheless have truth conditions because they exhibit conceptual contents which can be either true or false. The desire that it rain exhibits the content *that it's raining*, which is either true or false. Thus, in the case of cognitive representations, states' conceptual contents are their veridicality

[24] Veridicality conditions are sometimes called satisfaction conditions (e.g., Chalmers 2004, pp. 341-342).

conditions. If one desires that it rain, and it is raining, then the content of one's desire is true and one's desire is thereby fulfilled.

If the target of a representation is something like a state of affairs, then it is reasonable to hold that the representation has truth conditions. We typically describe conceptual contents in general with sentential clauses such as 'that a is F' (e.g., *that the tomato is red*). Because conceptual contents exhibit this sentence-sized character, they, like sentences in natural languages, can be true or false.[25]

But it is arguable that some representations do not represent there to be things like states of affairs. If a state represents there to be, for example, objects or properties, it is not obvious that those states could be the bearers of truth or falsity, though there may be conditions under which those representations are accurate. For example, suppose that a state represents not the state of affairs that the dog barks, but instead only the dog. The latter state cannot be true or false, but it is arguably accurate if the dog is present and inaccurate otherwise. Several theorists

[25] I note that there is some debate about whether all conceptual contents are sentential in this way. For example, using 'propositional' to mean having sentence-sized character and therefore being evaluable for truth or falsity, Crane observes that "some philosophers think that all representational content is propositional. Yet it seems to me that there are many intentional states whose contents are not assessable as true or false; for example, the object-directed attitudes of love and hate" (2009a, p. 478). We often say, for example, that Jack fears Jill. Here Crane presumes that Jack's fear is a conceptual state that represents Jill. But it would seem that the state's content does not exhibit sentential character and so does not exhibit a truth value. This claim is, however, controversial. These attributions of states involve so-called intensional transitive verbs, and there is a lot of debate about how such verbs operate. For example, Graeme Forbes (2006) argues that constructions involving such verbs are elliptical for constructions that do involve sentential that-clause compliments. So, for example, 'Jack fears Jill' is elliptical for another sentence, such as 'Jack fears that Jill is nearby', where the that-clause compliment 'Jill is nearby' does exhibit a truth value. I cannot settle these issues here. But whether or not Crane is correct that such constructions involve attributions of conceptual contents to states that are not truth evaluable, I will argue that mental qualities cannot be true or false, nor are they conceptual in any way.

have thus argued that veridicality conditions come in at least two varieties: Truth conditions and what I'll call accuracy conditions (e.g., Burge 2010, p. 9).

As Susanna Siegel notes, a state's exhibiting accuracy conditions involves comparing it to something else to identify whether it is accurate. Siegel writes,

> If something is accurate, then there is something else in relation to which it is accurate. If a map is accurate, then there is some spatial area in relation to which it is accurate. If a story about Simone is accurate in some respects but not others, then Simone has some of the features attributed to her by the story but not others. When a map, a story, a mental state, or anything else is accurate, there is a situation of which it is accurate. Attributing accuracy to something thus involves assessing it with respect to something else (2010, p. 31).

Consider, for example, a painting of George Washington. While we would not describe the painting as true or false, the painting can be accurate or inaccurate. For each property that the painting depicts George Washington as having that Washington did exhibit, the painting is accurate. For example, if the painting depicts Washington as having gray hair, and Washington did have gray hair, then the painting is accurate in that respect. If, by contrast, the painting depicts Washington as having brown eyes when he in fact had blue eyes, then the painting is inaccurate in that regard. Just as a cognitive state's conceptual content is true just in case the target state of affairs that the state represents there to be obtains, a mental representation that represents a property as obtaining is accurate just in case the property it represents there to be obtains.[26]

As I'll discuss shortly, there is a question about what explains why nonmental things such as paintings represent what they do—and even whether they are genuine representations at all.

[26] The accuracy of some representations such as paintings arguably comes in degrees. For instance, a painting of George Washington may be more or less accurate depending on how many properties it represents Washington as having that Washington did exhibit. A representation that only represents a single property cannot be more or less accurate—it is either accurate or inaccurate.

For example, why does a painting of Washington represent Washington as opposed to a collection of shapes or something else altogether? But whether or not nonmental things are genuinely representational, it has become increasingly common to argue that sensations and perceptions have accuracy conditions. Siegel writes,

> [This view] can be refined into a proposal that finds the following similarity between visual experiences and beliefs: like beliefs, maps, and newspapers, visual experiences have contents, and just as the contents of beliefs are conditions under which the belief state is true, so the contents of experiences are conditions under which the experience is accurate (2010, p. 30).

Sensations and perceptions thus arguably have accuracy conditions just as beliefs have truth conditions. In chapter 5, I will argue that these states have accuracy conditions in virtue of exhibiting mental qualities.

Before closing this section, I should mention an important debate about the nature of representation, in order to anticipate another kind of reservation that one may have about the idea that mental qualities are representational. For various reasons, many theorists have maintained that the only genuinely representational things are cognitive states and that if anything nonmental purports to be representational, it is only derivatively representational in virtue of being suitably related to cognitive states. John Searle (1992, pp. 78-79) calls this distinction the difference between intrinsic and derived intentionality (see, e.g., Haugeland (1998, pp. 127-170) for a similar distinction between original and derivative intentionality). According to Searle, cognitive states such as conscious thoughts exhibit representational content without having to be related to anything else (e.g., Searle 1992, p. 79).

On this view of intentionality, if anything other than a cognitive state is representational in any way, it must be due to being related to a cognitive state, which is genuinely and intrinsically representational. A book or a painting, for instance, may be representational in

some derivative sense. A robot may be described as having intentional states in some way. But, Searle maintains, these things exhibit only derived representational properties, parasitic on the intrinsic representational character that cognitive mental states enjoy (e.g., Searle 1992, pp. 79-80).

Whether or not Searle is correct that the intentionality of mental states is intrinsic, it is tempting to think that the intentionality of mental states is somehow the primary form of representation, whereas the intentionality of nonmental items is somehow dependent on that mental intentionality. If this is so, then it is also tempting to think that if mental qualities are representational at all, they are at best derivatively representational. But this would be mistaken. To begin with, the distinction between original and derivative intentionality is dubious for several reasons which, for considerations of space, I cannot discuss here (see, e.g., Dennett 1990). Moreover, even if there were a distinction between genuinely representational properties and merely derivative ones, the question remains whether or not nonconceptual mental properties are also genuinely representational.

Despite the initial appeal of representationalism, however, many philosophers still deny that mental qualities are representational. Why? The answer, I think, has to do with how many have approached the topic of mental qualities.

§1.3 Intuitions and the Puzzle of Mental Qualities

Proponents of the traditional qualia view often argue against representationalism by offering putative counterexamples to it. It has become something of a favorite pastime among qualophiles to attempt to find examples of paired perceptions within a given sense modality that are identical in terms of the perceptible properties perceived, but nonetheless qualitatively

distinct from one another (e.g., Peacocke 1983, Block 2003, Burge 2003). If such cases can be found, then it would seem that those differences must be due to those states' nonrepresentational qualia. Consider, for instance, Paul Boghossian's and J. David Velleman's (1989, p. 94) well-known case of blurry vision. Boghossian and Velleman urge that blurry vision and normal vision can concern the same things, though they appear to be qualitatively distinct. It would seem, then, that the difference between these states is not a difference in the perceptible properties that one sees, but a difference in how one sees those things—that is, a difference in those states' raw qualia.[27]

Representationalists, however, typically have persuasive replies to these purported counterexamples. That is, representationalists are always able to argue that such cases do involve differences in the perceptible properties that are perceived (for a review of these kinds of replies, see, e.g., Tye 2000 or Tye 2003). Blurry vision, for example, arguably presents boundaries in less detail than nonblurry vision does (e.g., Tye 2003, pp. 17-20). Thus, these cases do not establish that we are ever aware of qualia. Like many, I find these representationalist replies to such cases convincing, but, for reasons of space, I will not review this kind of debate in this dissertation. Instead, I will explore what I think is the foundation of the qualophiles' approach—that is, what I think motivates qualophiles to search for putative counterexamples to representationalism in the first place.

It is commonplace in philosophy to hold that the starting point of an account of a phenomenon is to accommodate intuitions about it. And there are several relevant intuitions that many philosophers claim to possess regarding mental qualities. Perhaps the most prevalent

[27] Qualophiles' counterexamples to representationalism continue to grow more and more recherché. Recent examples have involved experimental studies involving the deployment of peripheral attention (e.g., Block 2010) and ordinary experiences involving complex set ups with mirrors (e.g., Millar 2011).

intuition—one which goes back as far as antiquity—is the idea that so-called undetectable quality inversion without misrepresentation ("quality inversion" for short) is possible or at least conceivable (e.g., Sextus Empiricus 2005, Locke 1700/1975, Shoemaker 1982, Block 1990).[28] According to this intuition, it is possible that two people can veridically perceive the same object and yet exhibit distinct mental qualities, but in a way that cannot be detected even in principle. For example, some hold that it is possible that two people can veridically perceive a blueberry and behave in indistinguishable ways with regard to it, while at the same time one person's perception exhibits mental blue whereas the other person's perception exhibits mental yellow.

There are, of course, many everyday ways in which one person's mental qualities can differ from another person's. Some individuals are color blind and their perceptions of blueberries may exhibit mental qualities other than mental blue. But the premise of quality inversion crucially differs from these ordinary phenomena. The intuition requires that the differences between people lie in principle beyond what we can discover. Color blindness, by contrast, is plainly detectable. According to the intuition, even if we understood the neural organization of each of the so-called inverted twins, this knowledge would not help us to discover which brain exhibits which mental quality. It seems that a unit of brain tissue could have given rise to one mental quality just as easily as to the other. If this kind of scenario is possible, then it would seem that mental qualities are clearly not representational.

A related intuition about mental qualities is the intuition that so-called philosophical zombies are possible (e.g., Kirk 1974, Chalmers 1996). Philosophical zombies are creatures

[28] At this point, I note that the claim that something is conceivable is arguably weaker than the claim that the thing is possible, though many philosophers have tried to draw a strong connection between what is conceivable and what is possible. Others, however, deny that conceivability is a good guide to possibility (for illuminating essays on the topic, see, e.g., the papers in Gendler and Hawthorne 2002). In this dissertation, I'll address only the possibility of certain scenarios, not their conceivability, even though I recognize that these are different things.

purportedly physically and functionally identical to normal human beings, with the exception that their mental states lack mental qualities. Notably, the claim that philosophical zombies are possible is an even stronger claim than the claim that quality inversion is possible. If quality inversion is possible, then we cannot know which mental qualities a given individual exhibits. But if philosophical zombies are possible, we cannot know whether or not an individual exhibits mental qualities at all. Moreover, if we can conceive of zombies that veridically perceive their environments, then again it would seem that mental qualities are not representational.

In addition to having ramifications for the representational character of mental qualities, these kinds of intuitions are often claimed to drive us to powerful conclusions about the underlying nature of mental qualities. In particular, many use these sorts of scenarios to argue that mental qualities are not physical and that physicalism—the view that everything is physical—is false. If we must respect these intuitions, it is not hard to see how such conclusions follow. If it is possible for there to be a creature that is physically and functionally identical to a human being who exhibits mental blue, but a creature who exhibits a different mental quality (such as mental yellow) or no mental quality at all, then mental qualities are not identical to anything physical or functional—and hence physicalism is false (see, e.g., Chalmers 1996, p. 123 for an argument to this effect).

These seemingly perplexing features of mental qualities motivate what Chalmers (1996) has called the hard problem of consciousness and what Joseph Levine (2001) calls the explanatory gap. According to Chalmers, the easy problems within the study of the mind include tasks such as explaining those functions of brain processes that underwrite qualitative mental states. What makes these problems comparatively easy is that we possess at least an inkling of how they might be resolved. Moreover, these are the problems that can be left to the empirical

sciences. The hard problem of consciousness, by contrast, concerns why a brain state has the mental qualities it does, as opposed to other qualities or even no qualities at all. This problem, it seems, may have no satisfying resolution.

Not only do many philosophers think there is a hard problem, but also many scientists do so as well. Here Ryota Kanai and Naotsugu Tsuchiya proffer a rather standard characterization of the hard problem, pitched in a science journal for a general scientific audience:

> Perhaps the most difficult biological question of all might be how and why electrochemical neuronal activity in the brain generates subjective conscious experience such as the redness of red or the painfulness of pain. Neuroscientists track how light impinging on the retina is transformed into electrical pulses (neuronal spikes), relayed through the visual thalamus to reach the visual cortex, and finally culminates in activity within speech-related areas causing us to say 'red'. But how such experience as the redness of red emerges from the processing of sensory information is utterly mysterious…. This is called the 'hard problem' of consciousness as coined by the philosopher David Chalmers (2012, p. R392).

In a similar vein, Levine claims that there is a seemingly unbridgeable gap between our understanding of physical entities such as brain states and our understanding of mental qualities—even if mental qualities are also physical. These concerns lie at the core of one of the classic problems in the history of philosophy, the so-called mind-body problem. Such concerns are so pervasive that many today, including many cognitive scientists, do not study mental qualities at all because they regard such an enterprise as doomed from the outset.

§1.4 Consciousness and the Representational Character of Mental Qualities

There are additional reasons why these kinds of intuitions undermine the idea that mental qualities are representational. Though these intuitions are often cited as reasons to doubt the truth of physicalism, there is an important consequence of such intuitions that is less frequently noted. If it is possible that someone else's perception of a blueberry might exhibit mental yellow

instead of mental blue—and if there were no way to determine which quality it does exhibit—then it would appear that the only genuine access we have to mental qualities is through our own subjective awareness of them in consciousness. Of course, it is plain that we often have first-person access to our own mental qualities in consciousness. But if quality inversion is possible, then it follows not only that only I can know which mental qualities my own current conscious perceptions exhibit simply by examining my mind, but also that others cannot genuinely know them.

If, by contrast, we can have third-person, intersubjective access to one another's mental qualities, then we can in principle detect when others' mental qualities are present. In that case, quality inversion would be impossible. Thus, if scenarios such as inversion are genuinely possible, then it is in the nature of mental qualities to be available solely by way of consciousness.

In other words, the traditional qualia view depends on the contentious assumption that the only genuine access that we have to mental qualities is via our own subjective awareness of them in consciousness.[29] The approach to mental qualities echoes the Cartesian assumption that one's psychological states are always conscious and always fully known to one.[30] But this assumption takes Descartes a step further, holding that one not only always knows about one's own mental qualities, but also that subjective access is the only genuine access to them.

For most contemporary philosophers who accept the conceivability of quality inversion, this assumption operates simply as a backdrop to their theorizing. Some, however, explicitly

[29] Rosenthal (2010) has called views that make this assumption consciousness-based theories. The following discussion is much influenced by Rosenthal's treatment of these issues.
[30] For example, Descartes writes, "Now as to the doctrine that there can be nothing in the mind, insofar as it is a thinking thing, of which it is not aware, this appears to me self-evident" (2006, p. 147; AT VII, 246, Fourth Replies).

endorse it. Horgan and George Graham, for example, explicitly claim, "Phenomenal character, of course, is something that directly manifests itself only from within a first-person perspective on the world. What-it's-likeness is something one is acquainted with only in one's own case" (forthcoming, p. 338). Similarly, Kanai and Tsuchiya write, "The difficulty with qualia is their subjective nature: qualia exist only as viewed from the inside. They cannot be objectively detected or compared like any other properties measured in natural sciences" (2012, p. R392).

If how mental qualities appear in consciousness exhausts our knowledge of them, it follows that, when it comes to mental qualities, there is no appearance-reality gap. Speaking of the difference between appearance and reality in general, Nagel writes, "Experience itself, however, does not seem to fit this pattern. The idea of moving from appearance to reality seems to make no sense here" (1974, p. 444). In other words, a key consequence of traditional qualia theory is that it entails that how mental qualities appear in consciousness is precisely how they are.

The problem is that consciousness does not seem to give us much to work with in terms of explaining mental qualities. If one focuses on one's own perception of a blueberry, one can say that it exhibits some particular mental quality, but it appears there is not much else one can say. It certainly does not seem as though one is aware of an event in one's brain. Consciousness seems to leave us nothing with which to explain mental qualities or with which to identify them, either in the external world or within our brains. Thus Block rhetorically asks, "What is it that philosophers have called qualitative states? I answer, only half in jest: As Louis Armstrong said when asked what jazz is, 'If you got to ask, you ain't never gonna get to know'" (1978, p. 73). As Block claims, and as most other qualophiles would agree, we know about our own mental qualities via consciousness and we may not be able to say much more about them. Mental

qualities would seem to be simply mysterious epiphenomena, mental properties that play no psychological role.

If it appears from the first-person perspective that mental qualities cannot be explained in terms of anything else, then the traditional qualia view would have it that they must be inexplicable. This is why mental qualities are often characterized as, as Dennett puts it, those "properties of a subject's mental states that are 1. ineffable 2. intrinsic 3. private 4. directly or immediately apprehensible in consciousness" (1988, p. 229). And if mental qualities are intrinsic, then they are nonrepresentational. Thus, as Tye notes, many also characterize mental qualities as the "intrinsic, introspectively accessible, nonrepresentational qualities of experiences" (2002, p. 447). If mental qualities are not only nonrepresentational, but also nonphysical, intrinsic, ineffable, and so, then perhaps they are quite mysterious indeed.

§1.5 The Representationalist Approach to Mental Qualities

Proponents of the traditional qualia view often hold that mental qualities are by definition nonrepresentational. It is questionable, however, whether the above characterizations of mental qualities adequately capture our ordinary conception of them.[31]

Common sense holds that mental qualities are mental properties that we can become aware of from the first-person perspective in consciousness, properties which also connect to our thoughts and our speech in specific ways. Mental red, for example, is typically caused by sensing red things and typically causes one to think and say things such as "There is red." It is

[31] Cf. Kanai and Tsuchiya, who acknowledge that the "fact that we can conceive nonsensical possibilities such as philosophical zombies suggests that our conceptualization of qualia may still be premature" (2012, p. R393).

arguable, then, that our ordinary conception of mental qualities is simply silent regarding whether or not mental qualities are representational.

But what of the intuitions that have led many to conclude that mental qualities are nonrepresentational? As several representationalists have argued, it is questionable whether the scenarios these intuitions characterize are possible or even conceivable (e.g., Harman 1996a, p. 12; Tye 2002, p. 452). Moreover, as noted above, these intuitions depend on the critical assumption that mental qualities are only known via the first-person perspective. Many, of course, do not see this as a problematic assumption—or even as an assumption at all. As the above quote from Horgan and Graham shows, many simply think that there is no other route to investigate mental qualities.

If, however, there were another way to understand mental qualities—a way that does not assume that our only access to them is via consciousness—then the assumption underlying the traditional qualia theory is optional and arguably false. And there is such a different approach to mental qualities. According to representationalism, mental qualities are understood in terms of their roles in representing perceptible properties in perception.[32] This way of understanding mental qualities does not require that we have solely first-person access to them. In chapter 5, I'll argue that representationalism is to be preferred to theories that assume that consciousness provides the last word about mental qualities because representationalism not only accommodates our first-person access to mental qualities, but also captures our ordinary conception of mental qualities in a more adequate fashion than does the traditional qualia theory. Moreover, I'll argue that the intuitions that seem to support the traditional qualia theory are themselves unsupported. I'll then develop and defend a representationalist theory of mental

[32] Rosenthal (2010) calls theories that explain mental qualities in terms of their roles in perception perceptual-role theories. Representationalism is one kind of a perceptual-role theory.

qualities. With this theory in hand, we can see that scenarios such as quality inversion are actually impossible.

In the remainder of this chapter, I'll explore several other kinds of reasons that might motivate people to deny that mental qualities represent perceptible properties. I'll begin by briefly mentioning some alternative views, and then provide a roadmap of the various alternative views that I will challenge in the rest of the dissertation.

§1.6 Some Alternative Views and a Roadmap to the Dissertation

One motivation for resisting that mental qualities represent perceptible properties depends on one of the classic theories of representation. According to the so-called resemblance theory of representation, an item represents some target just in case that item suitably resembles that target (see, e.g., Goodman 1976, p. 3). On this view, a painting of Washington represents Washington because the painting suitably resembles Washington. But, as many have observed, particular mental qualities do not match, let alone resemble, particular perceptible properties, either characterized in commonsense color terms or in terms of their physical natures (cf. Rosenthal 2005, pp. 197-198). One might thus conclude that mental qualities cannot represent perceptible properties.

The resemblance theory has, however, been largely discredited for several reasons (see, e.g., Goodman 1976, p. 4). For example, most agree that the resemblance relation is symmetrical, but that the representation relation is not symmetrical. For example, if Washington

resembles Jefferson, then Jefferson resembles Washington. By contrast, if a painting of Washington represents Washington, it does not follow that Washington represents the painting.[33]

But even if one rejects the resemblance theory, there are several views about the relationship of mental qualities to perceptible properties that challenge the idea that the former represent the latter. Consider, for instance, the somewhat bizarre view of qualitative mentality that Aristotle proposes in De Anima. Though the nature of his view is a matter of some debate, Aristotle seems to have held that mental qualities are literally the same properties as the properties of objects (e.g., 1968, pp. 42-43; II.12 424a17-27; for discussion, see, e.g., Burnyeat 1992). On Aristotle's view, when one looks at a red tomato, one's eye literally turns red—that is, it acquires the very same property that the tomato exhibits. This view thus threatens the account that I'll develop because Aristotle identifies the qualities of perceptible objects with the qualities of mental states, whereas I will argue that these properties are distinct from one another and that the mental properties represent the perceptible ones. But Aristotle's view strikes most modern readers as hardly credible for many reasons, not the least of which is the fact that one's eye plainly does not turn red in the way that a tomato is red. Moreover, this view fails because perceptible properties are properties of three-dimensional objects, whereas mental qualities are properties of states of mind—that is, properties of states of objects (cf. Rosenthal 2005, pp. 196-197). Even if Aristotle were correct that one's sense organ takes on the same quality of what is perceived, it is difficult to see how a perception, which is a mental state, could acquire that perceptible property.

[33] Even if there were a kernel of truth in the resemblance theory, we need not explain how mental qualities represent perceptible properties in terms of an individual correspondence between each mental qualities and a matching perceptible property. It is open that there is a resemblance between mental qualities and perceptible properties at the level of the families of those properties—and I will pursue this route in chapter 5.

There are, however, other more widely-held views of mental qualities and perceptible properties that also threaten the view I will put forward. According to many historical figures such as Locke (1700/1975) and Hume (1739/2000), there is a distinction to be drawn between the so-called primary qualities and the secondary qualities (for contemporary versions of this view, see, e.g., Chisholm 1957, Boghossian and Velleman 1989). The primary qualities include properties such as shape and number. These properties, most would agree, are genuine properties of objects. By contrast, these theorists maintain that the secondary qualities, which include properties such as color and taste, are properties of mental states, which only seem to be properties of objects. Hume writes,

> The fundamental principle of [the modern] philosophy is the opinion concerning colours, sounds, tastes, smells, heat and cold; which it asserts to be nothing but impressions in the mind, deriv'd from the operation of external objects, and without any resemblance to the qualities of the objects (1739/2000, p. 149; §1.4.4).[34]

According to proponents of the primary-secondary quality distinction, though external objects may cause sensations of colors, those colors are not literally properties of the object. Proponents of the distinction are typically motivated by the fact that it seems from the first-person perspective that properties such as colors cannot be identical to physical properties such as the reflectance properties of the surfaces of objects. In other words, properties such as colors seem irreducibly qualitative. If so, then a natural conclusion is that properties such as colors reside only in the mind and somehow falsely appear to us to cover the surfaces of objects. There are, for example, mental colors, but there are no perceptible colors.

If there are no perceptible colors, then mental colors arguably do not represent them. If instead mental colors do represent there to be perceptible colors, then mental colors are always

[34] In more contemporary language, Boghossian and Velleman describe this view as the idea that "the representational content of visual experience represents external objects as possessing colour qualities that belong, in fact, only to regions of the visual field" (1989, p. 95).

misrepresentational. Either way, views that adhere to the primary-secondary quality distinction threaten the view that mental qualities often accurately represent there to be perceptible properties.

The primary-secondary quality distinction is, however, questionable. First, as Berkeley and others have observed, nothing is ever perceived to have a putatively primary quality such as shape if it is not also perceived to have some putatively secondary quality such as color.[35] The shape boundaries registered by vision, for example, are marked by color boundaries. It is thus unmotivated to regard objects as exhibiting properties such as shape while relegating properties such as colors to the mind because shape properties are perceived in tandem with, and partially in virtue of, the perception of colors.

Second, and perhaps more importantly, common sense does hold that perceptible properties are mind-independent properties of objects and it does not seem from the first-person perspective that perception involves the projection of qualities onto those objects. It would thus be preferable to accommodate perceptible properties if we can do so. I will therefore not argue further against this distinction per se, but I will defend a view that shows how we can accommodate perceptible properties, thus undercutting the motivation for the distinction.

Setting aside questions about the nature and existence of perceptible properties, there are also a variety of positions regarding the relationship of mental qualities to conceptual contents that would challenge the view that mental qualities are representational. For example, many historical figures, such as Aristotle (1968), Berkeley (1710/1996), and Hume (1739/2000), sought to construct conceptual content out of mental qualities (see, e.g., Prinz 2002 for a modern

[35] For example, Berkeley writes, "For my own part, I see evidently that it is not in my power to frame an idea of a body extended and moved, but I must withal give it some colour or other sensible quality which is acknowledged to exist only in the mind" (1710/1996, p. 28; Principles §10).

version of this position). For example, Hume famously maintained that "[a]ll ideas are borrow'd from preceding perceptions" (1739/2000, p. 399; Appendix). On this sort of account, the thought that there is a red tomato is somehow constituted by something like a faint version of mental red.

Since the view of mental qualities developed in chapter 5 sharply distinguishes mental qualities from concepts, it might seem that this kind of view threatens the view I will develop. But it is not obvious that these views are incompatible. In any case, I will not discuss this view in detail in this dissertation, though the view of mental qualities that I will develop does suggest reasons why this kind of view may not succeed: Since conceptual and qualitative representations are radically different kinds of properties, one cannot be built out of the other.[36]

In the course of developing this view, I will consider various alternative positions. As previously noted, most standard versions of representationalism depend on atomistic accounts of the contents of qualitative states. In chapter 2, I'll argue that such standard forms of representationalism fail primarily because of this atomistic approach to qualitative content.

For reasons enumerated above, standard forms of representationalism are incompatible with the possibility of undetectable quality inversion. Like these standard versions of representationalism, the version of representationalism that I will develop in chapter 5 denies that such inversion is possible. Recently, however, some representationalists have offered more sophisticated versions of the view to accommodate the possibility of inversion. But in chapter 3 I'll argue that these more complex accounts—including Shoemaker's dispositionalist

[36] Others deny that mental qualities are representational for different reasons. For instance, according to the view known as relational direct realism, perceptions are not representations at all; rather, perceptions are relations that hold between perceivers and perceived objects (e.g., Campbell 2002, Brewer 2006). I will not discuss this kind of view in detail in this dissertation, though my arguments for the representational character of mental qualities in chapter 5 will point towards why I find such positions problematic.

representationalism (e.g., Shoemaker 1994b) and Chalmers's Fregean representationalism (e.g., Chalmers 2004)—are unworkable.

One might think that the failures of these versions of representationalism suggest that the traditional qualia theory is true. But if the raw qualia view is correct, then we would require some explanation of how it is that we acquire knowledge of mental qualities, even from the first-person perspective. Perhaps the most prominent account of our knowledge of mental qualities is the theory that they function in our cognitive economies by being related to a special class of representations, so-called phenomenal concepts ("p-concepts") (e.g., Loar 1990/1997, Papineau 2002, Chalmers 2003, Block 2007b, Balog 2009). But I'll argue in chapter 4 that the traditional qualia theory is problematic because several versions of this p-concepts approach are unworkable.

I'll then argue in chapter 5 that the proper conclusion to draw from the failures of the accounts discussed in chapters 3 and 4 is that scenarios such as quality inversion are in fact impossible. I'll also argue that there are independent reasons to reject the intuitions behind these scenarios: These intuitions depend on several questionable assumptions, which are, on reflection, false.

In light of these considerations, I propose in chapter 5 a version of representationalism wherein qualitative content is individuated in a holistic way. This holistic theory of qualitative content—what I'm calling perceptual-role semantics—builds upon a burgeoning holistic theory of qualitative character, versions of which have been defended by, among others, Wilfrid Sellars (1963), David Lewis (1972), Sydney Shoemaker (1975, 1994a), Austen Clark (1993, 2000), and David Rosenthal (2005, 2010). On this holistic view of qualitative character, creatures' perceptual discriminatory capacities determine spaces of perceptible properties that reflect the

creatures' similarity-and-difference judgments regarding those properties. In order for creatures to generate such spaces, these theorist argue, creatures must be capable of being in states that vary in ways that mirror the ways in which the perceptible properties relate to one another. Thus mental qualities can be understood by their relative locations in spaces of mental qualities that match corresponding spaces of perceptible properties.

Though proponents of this view do not explicitly defend it as a theory of the contents of qualitative states, I build on the theory, arguing that the mental qualities are best understood as holistically individuated representations. For example, mental red represents perceptible red because it occupies a relative location within a space of mental colors that corresponds to the location occupied by perceptible red within the quality space of perceptible colors. The resultant holistic version of representationalism avoids the problems that plague other versions of it, resolves a host of philosophical puzzles about qualitative character, fits with a range of recent empirical findings about perception, and opens perception up to fruitful further study.

After explaining in detail how this kind of holistic qualitative representation differs from other kinds of mental representation, I'll then defend this account of qualitative representation from several objections and discuss how it clarifies a number of debates in the contemporary philosophy of mind. I conclude that this account of mental qualities enables us to do justice to our first-person impressions of them, to understand their roles in our mental lives, as well as to understand their relationship to the natural world.

CHAPTER 2: AGAINST CONCEPTUALIST AND ATOMISTIC REPRESENTATIONALISM

§2.1 Introduction

According to the view that I'll develop in chapter 5, mental qualities represent perceptible properties. In other words, I will develop a version of the view known as representationalism—the view that the qualitative character of a state is identical with (or supervenes on) that state's representational properties (e.g., Harman 1990, Dretske 1995, Tye 1995, Lycan 1996, Byrne 2001, Thau 2002, Jackson 2007, Crane 2009, Hill 2009, Speaks 2009, Pautz 2010). According to representationalism, for example, the reddish qualitative character of the perception of a ripe tomato simply is (or is determined by) the property of the state's representing the redness of the tomato.

As noted in chapter 1, there are many varieties of representationalism (for a survey, see, e.g., Chalmers 2004). In this chapter, I'll focus on the version of representationalism that Shoemaker has dubbed standard representationalism, which holds that "the phenomenal character of a color experience simply consists in its representing, its being "of," a certain array of colors in the scene being viewed" (2003, p. 255). Shoemaker attributes this view to, among others, Gilbert Harman (1990), Dretske (1995), Tye (1995), William Lycan (1996), and Alex Byrne (2001).

In holding that all mental properties are or depend on representational properties, representationalists maintain that perception and sensation as having something in common with thought. As Michael Martin observes, "[a]n intentional theory of perception claims that visual experiences have an intentional content that represents the world as being some way. This is to

see experiences as akin to propositional attitudes such as beliefs" (1994, p. 745).[37] As a result, representationalism is often read as rejecting the idea that perceptions or sensations themselves exhibit qualitative properties. Discussing the qualities that figure in perception, Tye writes,

> Tradition has it that these qualities are qualities of the experiences. Tradition is wrong. There are no such qualities of experiences…. Still there are qualities of which the subjects of visual experiences are directly aware in introspection. They are qualities of external surfaces (and volumes and films), if they are qualities of anything (2000, p. 49).

In other words, perceptions and sensations are qualitative, representationalists argue, only insofar as they exhibit suitable representational properties in virtue of which they represent perceptible properties.[38] Representationalism is thus often described as the view that "the qualitative character of a sensation is the property that sensation represents its representational object as having" (Levine 2003, p. 58).

Such a characterization of representationalism may strike one as paradoxical. How could a property of a mental state be identical to a property of the object that the state represents there to be? If I am aware of an tomato, for example, no property of my mental state is a property of the tomato. David Chalmers articulates this worry thus:

> Some representationalists, such as Dretske and Tye, put their view by saying that phenomenal properties are identical to certain represented external properties, such as physical redness. As I am putting things, that would be a category mistake. Phenomenal properties are by definition properties of subjects or of mental states, and physical redness is not (or need not be) such a property (2004, p. 342).

But this is simply a terminological issue. We can understand representationalism to be the view that the qualities of mental states simply are (or are determined by) the representations of the

[37] Indeed, as I'll discuss shortly, some representationalists maintain that perception is not just akin to, but is identical with, belief. But not all versions of representationalism are committed to this view. See §2.4 below.

[38] It is plain that we can have thoughts that represent perceptible properties without those thoughts' being perceptions. So there is a question about what representational properties, according to representationalism, distinguish thought from perception. I will address this question in §2.4 and in chapter 5.

qualities of objects that we represent there to be. And as Chalmers (2004, p. 343) himself notes, once we construe representationalism this way, it is evident that representationalism is coherent.

Even within the rubric of standard representationalism, there are many varieties of the view. In chapter 5, I'll develop a version of standard representationalism, but the account that I'll develop differs from standard varieties of standard representationalism. As noted in chapter 1, there is much disagreement about the nature of the representational content that figures in states such as perceptions and sensations. Many representationalists maintain that qualitative states exhibit so-called nonconceptual content, though such states may also exhibit conceptual content (e.g., Dretske 1995, Tye 1995, Crane 2009a). We can call varieties of representationalism that hold that qualitative contents are nonconceptual versions of nonconceptualist representationalism. But many representationalists do not specify whether the contents of qualitative states are conceptual or nonconceptual, and some are explicit that they are conceptual (e.g., Harman 1990, Byrne 2001, Pautz 2010a, Mandik 2012). Call this alternative view conceptualist representationalism. In chapter 5, I will develop a theory of qualitative character according to which mental qualities are a certain kind of nonconceptual content—that is, I will defend a version of nonconceptualist representationalism. Thus one of my central goals in this chapter is to argue that qualitative content is nonconceptual.

But whether or not qualitative content is nonconceptual, representationalists must, of course, provide a credible account of how qualitative representation works. Representationalists, if they provide an account of the contents of perceptions, often endorse a version of an atomistic theory of content (e.g., Dretske 1995, Tye 1995, Lycan 1996, Hill 2009; for an overview, see, e.g., Neander 2006). On this kind of view, a state's representing a thing is a matter of that state's tracking—that is, detecting, carrying information about, causally co-varying with, or standing in

some other tracking relation to—that thing. I'll call versions of representationalism that depend on such tracking psychosemantics versions of atomistic representationalism (for a survey of versions of the view, see, e.g., Bourget and Mendelovici forthcoming).[39] On this view, for example, a perception exhibits a reddish qualitative character because it tracks red. In chapter 5, I will instead offer a theory according to which mental qualities represent perceptible properties in a holistic way. Thus my other central goal in this chapter is to demonstrate that such atomistic versions of representationalism are problematic.

Before providing reasons to think that qualitative content is nonconceptual, I will begin by considering what is perhaps the most promising argument for conceptualist representationalism.

§2.2 The Arguments from Justification and Grounding

Qualitative states such as perceptions are often thought to provide reasons or justifications for certain thoughts. If I perceive something red—and thereby exhibit mental red—my perception justifies me in believing that there is something red. I'll call thoughts that concern what it is that one currently perceives perceptual thoughts. And several theorists, most notably John McDowell (1994), argue that it would seem that perceptions can justify perceptual thoughts only if both perceptions and thoughts exhibit the same kind of content—namely, conceptual content. McDowell writes that "we cannot really understand the relations [to perceptions] in virtue of which a judgment is warranted except as relations within the space of concepts…" (1994, p. 7).

McDowell's claim here echoes Donald Davidson's oft-cited dictum that "nothing can count as a reason for holding a belief except another belief" (1986, p. 310). While McDowell

[39] This view is often called tracking representationalism (see, e.g., Cutter and Tye 2011),

does not endorse this very strong claim, he does endorse the weaker claim that, in order for states to stand in appropriate justificatory relations, they must both have the structure of thought—that is, they must both exhibit conceptual content (see, also, e.g., Brewer 1999). Since McDowell holds that perception justifies thoughts, he concludes that those states must exhibit conceptual content. We can call this the Argument from Justification.

Before discussing this argument, I should mention that there is a similar argument in the vicinity as well. Pautz (2010a, p. 266) proposes yet another sort of argument for the view that perceptions exhibit conceptual content, which I'll call the Argument from Grounding. Pautz claims that there is a pervasive intuition about perception, which he calls the grounding intuition. Describing what he calls H, a hallucination of something's being red and round, Pautz glosses this intuition as follows:

> [I]ntuitively, it is necessary that, if an individual who has the capacity to have belief at all has H (for a sufficient period of time), then he will thereby have the additional capacity to have a general belief that is true only if something or other is present that is red and a general belief that is true only if something or other is present that is round (2010a, p. 266).

Pautz claims that some creatures may not have the capacity to form perceptual thoughts, though they may be able to perceive. So he hastens to urge that his claim is only that, were a creature to have the capacity to form perceptual thoughts, certain perceptions ground that capacity. I register at this point that it is not entirely clear what Pautz thinks it is for some capacity (such as the capacity to perceive) to ground another capacity (such as the capacity to form perceptual thoughts). But I assume that Pautz conceives of grounding as some sort of explanatory relation. That is, Pautz holds that a creature's capacity to perceive wholly or at least in part explains the creature's capacity to form perceptual thoughts.

The grounding intuition figures in an argument for the view that perception exhibits conceptual content, which Pautz claims takes the form of an inference to the best explanation. Pautz argues that the best explanation of the capacity of perceptions to ground perceptual thoughts is that both states are representational states with the same kind of content. Pautz writes,

> [T]here is an argument to be made for the representational view of visual experience. It takes the form of an inference to the best explanation. Both veridical and nonveridical visual experiences can ground the capacity to have beliefs about the external world…. The best explanation of these and other features of visual experience, I argue, is that both veridical and nonveridical experiences are themselves representational states [with the same kind of content] (2010a, pp. 254-255).

The Arguments from Justification and Grounding are similar, but, as Pautz (2010a, p. 267) notes, distinct. We arguably cannot explain how perceptions justify perceptual thoughts unless we first explain how perception grounds thought, whether or not those perceptual thoughts are justified. In any case, both arguments purport to establish the conclusion that perception exhibits conceptual content.

It is worth noting at this point, however, that neither of these arguments entails that nonrepresentational qualia do not exist, nor do these arguments entail even the claim that such qualia, if they exist, supervene on representational content. Both arguments are consistent with the view that perceptions exhibit not only conceptual contents which ground or justify certain thoughts but also exhibit nonrepresentational qualia that do not supervene on content. Nevertheless, these arguments threaten the version of representationalism that I'll develop in chapter 5. Pautz's inference to the best explanation arguably relies on the following implicit assumption: Since conceptual properties alone are sufficient to explain how perception grounds thought, there is no explanatory benefit to positing nonconceptual contents or nonrepresentational qualia. Although McDowell does not explicitly avow such an inference, it

too seems to be present in his thinking.[40] Since mental qualities would be epistemically unnecessary, it would seem there is no reason to posit them. And even if mental qualities did exist, these arguments suggest that they need not be representational because it is qualitative states' conceptual contents, and not their qualities, that ground and justify our perceptual thoughts.

But both the Arguments from Grounding and Justification arguably fail to establish qualitative states exhibit conceptual content, and both fail for the same reasons. These arguments succeed only given the assumption that the only properties in virtue of which perception can justify or ground thoughts are conceptual contents. If nonconceptual contents can perform these grounding and justificatory functions too, then both arguments beg the question in favor of a conceptualist view.

There are, however, some reasons to think that nonconceptual contents cannot justify thoughts. One putative reason for this conclusion is that if nonconceptual representations are not evaluable for truth or falsity, then nonconceptual representations cannot stand in logical relations to thoughts. As McDowell would say, adapting Wilfrid Sellars's phrase, if mental qualities are not conceptual, they cannot stand in "the logical space of reasons" (e.g., McDowell 1994, p. xviii). If justificatory relations are a species of logical relations, then nonconceptual states can at best cause corresponding perceptual thoughts, but cannot justify those thoughts. As Davidson claims,

[40] For example, McDowell writes, speaking of Sellars's view that there are nonconceptual sensory aspects to experience, "And it is not clear why it should seem necessary to describe these suitably similar impacts in terms of nonconceptual impingements on consciousness (sensations), as opposed to saying that consciousness comes into play only with conceptual episodes, triggered by nonmentalistically described impacts on sensory equipment. It seems that what Sellars here introduces as proximate causes of sensations can themselves meet the explanatory need…. The sensations look like idle wheels" (1998, pp. 443-444).

> The relation between a sensation and a belief cannot be logical, since sensations are not beliefs or other propositional attitudes. What then is the relation? The answer, I think, is obvious: the relation is causal. Sensations cause some beliefs and in this sense are the basis or grounds of those beliefs. But a causal explanation of a belief does not show how or why the belief is justified (1986, p. 311).

Another major motivation for this view of justification is the idea that, if a perception or sensation is to provide a reason for a thought, one must be able to think about the relationship between the two states and cite the qualitative state as one's reason for one's thought. For example, McDowell writes,

> If these relations [between sensation or perception and the beliefs those states give us reason to form] are to be genuinely recognizable as reason-constituting, we cannot confine spontaneity within a boundary across which the relations are supposed to hold. The relations themselves must be able to come under the self-scrutiny of active thinking (1994, p. 53).

On McDowell's view, in order for the relationship between qualitative state and thought to be justificatory, one must be able to scrutinize or understand the rational transition between that state and the thought. And, in order to understand the transition between these states in this way, it would seem that one must be able to have thoughts that exhibit the same contents as both of those states. But if qualitative states do not exhibit conceptual contents, then one cannot have thoughts that exhibit those contents. So, on this conception of justification, qualitative states cannot function as reasons for perceptual thoughts.

Despite the pervasiveness of this picture of justification among contemporary epistemologists, the idea that there are nonconceptual representations that can justify thoughts is also common among contemporary philosophers of mind. As several proponents of nonconceptual content have observed, the fact that a nonconceptual state is representational is arguably sufficient for those states to function as reasons (e.g., Heck 2000, pp. 504-506). But what about the putative requirement that a state can only function as a reason if one is capable of

having a thought with the same content as that state? As several proponents of nonconceptual content have argued, this requirement is far too strong. Instead, these proponents of nonconceptual content have argued that one need only be able to think about one's nonconceptual state for it to function as a reason for thought. For example, Jeff Speaks writes,

> [T]he nonconceptualist can admit that if something is a reason for a subject, the subject can talk about it; the nonconceptualist need only deny that if some proposition or content is a reason for a subject to do something, that subject must be able to utter some sentence that has that content as its meaning (2005, p. 395, fn. 21).

Speaks' point here is that, even if sensations cannot figure in thoughts or even if one cannot utter a sentence that has the same content as the sensation's content, one can nonetheless think about one's sensations and thus they can function as reasons for one's thoughts. Thus the requirement that one be able to cite one's states as reasons for them to count as reasons may be too stringent. Infants' or nonhuman animals' qualitative states arguably justify their perceptual thoughts, even though they cannot cite their those states as reasons. After all, what causes one to form a perceptual thought is not one's thought about its corresponding qualitative state, but the qualitative state itself. In causing the perceptual thought, the qualitative state is the reason for the thought. I cannot provide here a complete account of what it is that makes a state count as a reason. But whatever the conditions on something's being a reason might be, it is possible that mental qualities can function as reasons even if they are not conceptual.

The idea that justification holds only between representations of a specific kind depends on a demanding Davidsonian conception of justification, according to which justification is a logical relation holding between conceptual states. But the Davidsonian conception is questionable and not the only way to understand justification or epistemology in general. On a Quinean sort of naturalized epistemology, by contrast, epistemology centrally concerns the question of how creatures manage to navigate the world and form complex representations of it

given the comparatively meager evidence of their senses (e.g., Quine 1969). This question, Quine maintains, is a question for empirical psychology.[41] A naturalized theory of justification thus concerns the scientific explanation for why creatures make certain transitions between sensation and thought. On this sort of picture, mental qualities arguably can justify thoughts.

One could, however, dig in one's heels and insist that justificatory relations only hold between conceptual states that can enter into inferential relations, but it is not clear why one should do so. And even if it were the case that only a conceptual state can justify another conceptual state, it is an open question whether there is nonetheless a genuine epistemic relationship akin to justification that holds between nonconceptual mental qualities and conceptual thoughts. So there are no clear obstacles to thinking that mental qualities cannot justify perceptual thoughts, even if those mental qualities are nonconceptual. In other words, the Arguments from Justification and Grounding do not establish conceptualist representationalism, but they do establish at a minimum that qualitative states exhibit content. In chapter 5, I'll argue that mental qualities are not conceptual and that there are good reasons to think that they are the contents that play these roles in grounding and justifying perceptual thoughts.

§2.3 Familiar Arguments for Nonconceptual Content

I am not, of course, the only one to argue that sensation and perception involves representations that are not of a conceptual variety. Many theorists have argued that, though thought exhibits conceptual content, qualitative states exhibit nonconceptual content in addition to whatever conceptual content they may exhibit (e.g., Evans 1982, Peacocke 1992, Dretske 1995, Tye 1995,

[41] Quine writes, "The stimulation of his sensory receptors is all the evidence anybody has had to go on, ultimately, in arriving at his picture of the world. Why not just see how this construction really proceeds? Why not settle for psychology?" (1969, p. 75).

Bermúdez 1998, Fodor 2007, Crane 2009a, Burge 2010). There are a variety of accounts of nonconceptual content on offer, some of which I will discuss shortly.[42] But whatever the account of nonconceptual content, there are many reasons to think that qualitative states exhibit such content.

Robert Hanna and Monima Chadha provide a list of "at least seven different arguments for Non-Conceptualism" which they cull from "the recent and contemporary literature on mental content" (2011, p. 188; for a similar list of arguments, see, e.g., Speaks 2005, pp. 362-363). In what follows, I'll briefly mention a few of these arguments.[43] Perhaps the most commonly-cited argument for nonconceptual content is what I'll call the Argument from Fineness of Grain (Hanna and Chadha's first argument). According to this argument, perception exhibits nonconceptual aspects because we can perceive many various perceptible properties in a very fine-grained way, which arguably outstrips the creature's repertoire of comparatively course-grained concepts. Christopher Peacocke glosses this argument thus:

> There are many dimensions—hue, shape, size, direction—such that any value on that dimension may enter the fine-grained content of an experience. In particular, an experience is not restricted in its range of possible contents to those points or ranges picked out by the concepts—red, square, straight ahead—possessed by the perceiver (1992, p. 68, emphasis his).

[42] As Paul Coates (2007, p. 12) observes, it is unclear whether some uses of 'nonconceptual content' refer not to the contents of representations that are not conceptual, but rather to low-level conceptual contents—that is, contents that involve simple concepts such as concepts that pertain to perceptible properties. Cf. Sellars, who writes: "I also think, however, that most influential accounts mistakenly assimilate impressions, thus understood, to minimal conceptual representations, minimal both with respect to the physical content in terms of which they are described, and in being bracketed by looks or seems" (1967, p. 15). Coates suspects that the putatively nonconceptual contents that, for example, Christopher Peacocke (1992) posits may be low-level conceptual contents. I will discuss Peacocke's account shortly and I will offer other criticisms of it. But if some representationalists' nonconceptual contents are simply conceptual contents of a simple variety, then those views are open to the sorts of criticisms I will level against conceptualist representationalism in this chapter.

[43] For reasons of space, I will not consider all of these arguments.

The origin of this argument traces back to a point famously made by Gareth Evans, who rhetorically asks: "Do we really understand the proposal that we have as many color concepts as there are shades of color that we can sensibly discriminate?" (1982, p. 229). For example, it does seem as we can perceptually discriminate a huge range of perceptible colors, even thought it does not seem as though we possess concepts for all of those colors. If so, then it would seem that there are perceptual representations of these properties which are not conceptual (for similar versions of this argument, see, e.g., Tye 2006).

I mention the Argument from Fineness of Grain because it is very commonly cited as a justification for nonconceptual content, even though I think it fails to establish that perception exhibits representational aspects that are not conceptual. Peacocke is doubtless correct that we do not possess noncomparative concepts, concepts such as RED or BLUE, for every color that we can perceptually discriminate. However, the argument is unconvincing because we arguably do possess sufficiently many comparative concepts, such as A SHADE OF RED MORE ORANGE THAN RED or A SHADE OF BLUE DARKER THAN THE SHADE TO THE RIGHT to characterize every color we can consciously discriminate (cf. Rosenthal 2005, pp. 188-189; Mandik 2012, p. 619). Of course, that fact alone does not establish that perception involves solely conceptual content. After all, it may be that we possess concepts for every color we can perceive and yet do not deploy those concepts in perception (cf. Bermúdez 2007, p. 60; Berger 2012, p. 634). However, the fact that we might possess sufficiently many concepts for everything we perceive undercuts the Argument from Fineness of Grain.

A more persuasive argument for nonconceptual content is what I'll call the Argument from Infant and Nonhuman-Animal Perception (Hanna and Chadha's third argument). Peacocke offers the following gloss of this argument:

> The most primitive aspects of representational content in perception, which our subjective experience shares with the mere animals, do not involve the grasp of objectivity required for conceptual content. This is one of the reasons that trying to treat all qualitative content as conceptual involves an overascription (2001, p. 264).

According to Peacocke, infants and nonhuman animals interact in complex ways with the perceptible properties in their environments, which is evidence that they are capable of perceptually representing those perceptible properties. But it would seem that infants and nonhuman animals do not or need not possess concepts in order to perceptually represent those properties. Thus it would seem that infant and nonhuman animal perception involves nonconceptual representations (for similar versions of this argument, see, e.g., Dretske 1995, p. 11).

In a closely-related vein, many posit nonconceptual content to explain how it is that creatures acquire concepts via perception in the first place. Call this the Argument from Concept Acquisition (Hanna and Chadha's sixth argument). Tye glosses this argument thus:

> Another motivation for the nonconceptualist thesis is provided by the phenomenon of concept acquisition. Human beings acquire many general concepts via perceptual experience. How is this possible? If the world pre-conceptually really were a "blooming, buzzing confusion", as William James supposed, then it is utterly mysterious how we could acquire color and shape concepts from our experience. By contrast, if the world is already presented to us, at a nonconceptual level, as being made up of three-dimensional surfaces of varying colors and shapes, it is easy to understand how we can come to conceptualize the world via color and shape concepts (2005, p. 228).

If we assume that creatures are not born with concepts, then we require some mechanism to explain how it is that they acquire concepts. The most natural explanation is that we acquire concepts via our sensory encounters with the world. And many thus conclude that perception involves nonconceptual representational aspects.

Perhaps the best reason to think that qualitative states exhibit content is not explicitly cited by Hanna and Chadha, though it is related to the Argument from Concept Acquisition.

This argument turns the Arguments from Justification and Grounding for conceptualism, as it were, on their heads. Conceptualists argue that our perceptual thoughts could be justified or grounded only if perception and sensation are conceptual. But the nonconceptualist can retort that the fact that we have appropriate perceptual thoughts in the presence of perceptible properties in the first place at all calls out for explanation. Since concepts are comparatively complex mental representations, what could explain the presence of these complex states, if not that there are caused by simpler nonconceptual representations keyed to our environments? Thus many conclude that the close tie between sensation and perceptual thought is best explained by the fact that qualitative states are nonconceptual. Call this the Argument from the Connection of Sensation and Perceptual Thought.[44]

Unlike the previous two arguments, this argument does not provide reason to think that perception exhibits representational properties that are distinct from and somehow more basic than conceptual representations. It could be, for example, that sensation exhibits conceptual content, which explains its close ties to perceptual thoughts. But taken together, the above arguments do suggest that perception involves nonconceptual representations.

But there is a different problem with the forgoing arguments. Though these arguments point towards a version of nonconceptualism, they fail to establish a version of strong representationalism. That is, the arguments at best support a version of weak representationalism, according to which qualitative states exhibit nonconceptual content,

[44] Cf. Sellars's so-called sense-impression inference: The inference that there must be nonrepresentational sensory representations to explain these presences of these sorts of thoughts. Sellars writes that "the sense impression inference is an attempt to account for the fact that normal perceivers have conceptual representations of a red and rectangular object both (a) when they are being affected in normal circumstances by a red and rectangular object; and (b) when they are being affected in abnormal circumstances by objects which have other, but systematically related, characteristics" (1967, p. 17, emphasis his).

although they may also exhibit nonrepresentational mental qualities. And some theorists do endorse a view like this. For example, Peacocke (1992) himself clearly believes that perception exhibits nonconceptual content, but he is explicit that perception also exhibits nonrepresentational qualia.

My approach in what follows will be two-fold. My goal in this chapter is modest: To show that there must be a nonconceptual element to perception and sensation. Then, in the remainder of this dissertation, I will argue that the traditional qualia view is not only unmotivated but also problematic for other reasons. For example, I will discuss Peacocke's view in detail in chapter 5, wherein I will argue that positing nonrepresentational qualia in addition to nonconceptual content posits too much. My argument is in that way abductive. Taken together with the arguments for the view that there must be a nonconceptual element to qualitative states, a version of strong representationalism will be quite attractive.

I note that I have only briefly surveyed these familiar arguments for a nonconceptual element to qualitative states. But before moving away from conceptualist representationalism, I'll discuss in the next section what I think is the best reason to reject that view.

§2.4 The Best Argument against Conceptualist Representationalism

I think the best reason to doubt conceptualist representationalism is an argument that I'll call the Argument from the Distinction between Perception and Thought (Hanna and Chadha's fourth argument). The core of this argument is that a version of representationalism that denies that there are nonconceptual contents cannot explain the difference between perceptions and thoughts. Folk psychology plainly recognizes this difference: For example, there is a clear difference between the thought that there is something red on the table and the perception that

there is something red on the table. As noted in chapter 1, representational states are typically individuated by their contents and their mental attitudes, so conceptualist representationalists could locate the difference between thoughts and perceptions as differences in those contents, attitudes, or some combination of the two. Perhaps the most frequent way that representationalists distinguish perception from thought is through the proposal that perceptions exhibit a fundamentally distinct sort of content than thoughts—and I'll develop a version of this view in chapter 5. But I'll argue that if one denies that either nonconceptual contents or nonrepresentational qualia exist, then one cannot explain this difference.

Another common way that representationalists have distinguished thought from perception is to regard perceptions as exhibiting what can be called sensory or perceptual attitudes towards their contents. Susanna Schellenberg describes this proposal thus:

> This version of the [representationalist] thesis posits… that experience is a matter of standing in a certain attitudinal relation to [an representational content], analogous to the sense in which one might say that belief is a matter of standing in the believing relation to the [representational] content of the belief. English does not have a word to denote such a perceptual attitudinal relation. Byrne (2009b, p. 437) calls the relation the ex-ing relation; Pautz (2010a, p. 54 [*sic*, p. 258]) calls it the sensorily entertaining relation; Siegel (2010, p. 22) calls it the A-relation (2011, p. 730).

Representationalists such as Byrne and Pautz unfortunately do not provide much explanation of these perceptual attitudes. Pautz does explicitly mention that perceptual attitudes are distinct from other psychological attitudes as recognized by folk psychology, such as belief and desire. Pautz writes, for example, that "veridical and nonveridical experiences are themselves representational states of a kind more basic than belief" (2010a, p. 255).[45] And as Byrne observes (2009b, p. 450), these attitudes are not factive. Perception can clearly represent the world to be ways that it is not. And whatever these perceptual attitudes are, they are assertoric.

[45] On Pautz's view, experiences are more basic than beliefs insofar as they ground the capacity for thought.

Just as a belief that it's raining represents the world as being a certain way and therefore asserts the content *that it's raining*, a perception also represents the world to be a certain way. So if perceptions exhibit content, they assert those contents, though those contents need not be veridical. As I'll argue shortly, this account as it stands does not constitute a viable account of perception and representationalists cannot offer a further informative account.

Among those representationalists who claim that there are perceptual attitudes, there is disagreement regarding whether there are several distinct perceptual attitudes for each sensory modality or one kind of perceptual attitude. For example, though Dretske holds that perception exhibits nonconceptual content, he argues that there is a single perceptual attitude, which he calls the property-awareness relation (1999, p. 106), and that all differences between perceptual states are differences in the contents that one is related to via that relation. Considering the example of the difference between seeing something move and feeling that thing move through the sense of touch, Dretske writes,

> For there is more—much more—involved in seeing an object move than experiencing the object's movement. One also experiences the object's shape, size, color, direction of movement, and a host of other properties. This is why seeing and feeling movement are much different even though the same thing (movement) is represented in both modalities. Even when the senses overlap in their representational efforts—as they do in the case of spatial properties—they (McGinn 1991, p. 35) represent different ranges of determinable properties (1995, p. 95).

According to Dretske's version of representationalism, a visual perception of movement and a tactile perception of that movement both exhibit the property-awareness relation (what I am calling the perceptual attitude) towards contents involving movement. What distinguishes them is that the former exhibits content that involves, for example, the color of the thing that is moving, whereas the latter exhibits content that involves, for example, the temperature of the thing that is moving.

By contrast, some representationalists hold that there are distinct perceptual attitudes specific to each modality. Harman seems to endorse this sort of view when he writes:

> I am certainly inclined to think that there can be different attitudes with the same content and that what it is like to be in one of these attitudes need not be the same as what it is like to be in another. For example, what it is like to believe that P is not the same as what it is like to desire that P. Similarly, what it is like to see something as ahead and to the right is not normally the same as what it is like to hear something as ahead and to the right (1996b, p. 76).

On this view, a visual perception that there's a dog and a tactile perception that there's a dog can exhibit the very same content *that there's a dog*, but the difference between them consists in one state's exhibiting a visual perceptual attitude towards that content whereas the other exhibits a tactile perceptual attitude towards it.[46] Harman's version of representationalism does not imply therefore that all differences in perceptual states are differences in content. Rather, it holds that all differences in perceptual states are differences in either content or attitude. But since both content and attitude are representational properties, this view nonetheless counts as a version of representationalism.

The difference between these two sorts of representationalism does not matter for present purposes. Whether or not there is a single perceptual attitude or distinct perceptual attitudes for each of the sensory modalities, the representationalist owes an informative account of them. The proponent of perceptual attitudes might claim that they are explained functionally. Thus Lycan describes his version of representationalism as the view that there is "something to phenomenal character besides representational content. It might, in particular, be some aspect of functional

[46] Byrne calls the view that there is a single perceptual attitude and that differences between modalities are differences in content intermodal representationalism, which he contrasts with what he calls intramodal representationalism (2001, p. 205). As Byrne notes, Block (1996, pp. 37-38) draws a similar distinction between what he calls representationism and quasi-representationism. Byrne attributes intermodal representationalism to, among others, Dretske (1995) and Tye (1995) and intramodal representationalism to Harman (1996), Lycan (1996a), and others.

role," which Lycan claims is consistent with the view that "within a single [sense] modality, all phenomenal differences are representational differences" (1996a, p. 134). It is, however, incumbent upon the Lycan-style representationalist to say more about what these functional roles that define perceptual attitudes might be.

Part of the functional roles of the mental attitudes of thoughts is that they correspond to the illocutionary forces of the speech acts that express them (e.g., Austin 1962b, Vendler 1972, Searle 1983). If I, for example, say "It's raining," my speech act expresses my thought that it's raining. And my speech act does so by matching my thought in two ways: My thought exhibits the representational content *that it's raining*, and my speech act matches this by exhibiting the speech act content *that it's raining*. Likewise, my thought exhibits the belief mental attitude and my speech act matches this by exhibiting an assertoric illocutionary force. I can, of course, express thoughts with the same content but with different attitudes by making other speech acts. For example, if I wonder whether it will rain, I can express that state of wonder by asking "Will it rain?" Here, the speech act again exhibits the speech act content *that it's raining* (which matches the wondering's corresponding representational content), but it exhibits an interrogative illocutionary force to match the state of wondering's attitude.

Thus each mental attitude corresponds to a distinctive illocutionary force.[47] But it is unclear what distinctive illocutionary force could correspond to a perceptual attitude or attitudes. If I say "There's red," this is an expression of my thought that there is red. I might, of course, say "There's red" when I am perceiving something red, but what the representationalist needs is some reason to think that this speech act's illocutionary force matches the perceptual attitude, and not simply the attitude of a nonperceptual thought that accompanies the perceptual state.

[47] For an argument to this effect, see, e.g., Vendler 1972, chapter 3.

Indeed, there seems to be no reason to think this. Moreover, it is unclear what grip we have on mental attitudes if not in terms of the illocutionary forces of the speech acts that express them. Perceptual attitudes might be a special case insofar as they do not, unlike the other attitudes, correspond to distinctive illocutionary forces. But if so, then the representationalist owes another account of them.

Crane offers an analysis of a case, however, which seems to suggest that perceptual attitudes are not the kinds of things that correspond to illocutionary forces. Crane rejects the standard representationalist explanation of blurry vision discussed in chapter 1, according to which blurry vision represent things as having indistinct contours. Rather, Crane claims that the blurry vision exhibits the content that things have distinct contours; the blurriness is explained by the ways that one exhibits those contents, which Crane calls those states' representational modes (2009b, p. 483). If representational modes are mental attitudes, then it is not clear how the blurriness of vision could correspond to the illocutionary force of any speech act. Indeed, it is not clear what account Crane could give of such attitudes or how this view substantively differs from a view that posits mental qualities.

Crane notes that the main difference between this view and the view that blurry vision is explained by invoking nonrepresentational qualia is that it preserves the representationalist's thesis that all mental differences are or supervene on representational differences (2009b, p. 483). But without an informative account of these kinds of perceptual attitudes, the proposal is at best an effort to save representationalism. This result may not trouble Crane because, as he acknowledges, he doubts "whether representationalism can really help a reductive account of mind" (2009b, p. 487). But in the absence of an actual account of such attitudes, this proposal

merely re-labels the problem of explaining the difference between perception and thought and arguably invokes mental qualities without explaining them.[48]

It may, however, be that perceptions do exhibit one of the attitudes found in folk psychology. If so, it is unlikely that they exhibit, for example, the desire attitude because desire does not to commit one to the world's being a certain way, whereas perception does. The most natural suggestion is that the attitude is belief. Beliefs, after all, do commit one to the world's being some way, but they certainly can misrepresent the world as being a way that it is not. Describing the attitude involved in perception as the exing relation, Byrne also briefly considers the idea that this relation simply is belief. Byrne writes,

> Finally, the second question, about the nature of exing. It is supposed to be a non-factive propositional attitude that is constitutively involved in perception, but this is to say alarmingly little. If there really is such an attitude, it is puzzling that there is no corresponding propositional-attitude verb. These vexing issues would vanish if it turns out that exing is believing (2009b, p. 450).

Soon after articulating this position, however, Byrne shies away from endorsing the view (2009b, p. 451).[49] But the view that these states are simply kinds of belief has been defended recently by Kathrin Glüer (2009). She calls this view the doxastic account of perception.[50, 51] This view

[48] Chalmers similarly suggests that experiencing some content is "equivalent to the representational property of perceptually phenomenally representing the relevant… content" (2004, p. 365). But it is unclear what it is to perceptually phenomenally represent some content. If it just involves mental qualities, it would seem that this view smuggles in mental qualities without explaining them. I will address Chalmers view in more detail in chapter 3.

[49] Byrne (2009b, p. 451) then briefly considers the proposal that the perceptual attitude may involve, but not be identical to, belief. But Byrne does not explain what else perceptual attitudes might involve and thus fails to offer an informative account of them.

[50] The view that perception is a kind of belief often attributed to David Armstrong (e.g., 1968, p. 209). But, as Glüer herself points out (2009, p. 298, fn. 1), this is a misattribution. Armstrong did not hold that perceptions are beliefs, but rather that they are acts that involve the acquisitions of beliefs. There is then a remaining question what Armstrong's view of the nature of those perceptual acts might be, but I will not address that question here.

[51] There is another view in the vicinity that is worth mentioning. Many theorists have held that perception is a conceptual state that is appropriately caused by nonconceptual sensations. For

would seem to have the benefit of demystifying the attitude involved in perception. It would also neatly explain how it is that perception and belief interact—the answer is simple: They interact just as other beliefs do.

But if perception is a kind of belief, then the difference between perception and belief must be a difference in their contents. Glüer proposes that perceptions, unlike thoughts, represent perceiver-dependent properties that she calls phenomenal properties (2009, p. 311).[52] Without considering this proposal at length, I will observe only that it is controversial whether perception represents perceiver-dependent properties. In chapter 3, I will discuss a version of this proposal, Shoemaker's account of so-called appearance properties, and will argue that his proposal is problematic.[53] Then, in chapter 5, I'll argue that there are no such properties represented by perception. For now, I will assume that the differences in content must be differences in the typical viewer-independent properties represented.

It is worth noting, however, that even if perception did represent perceiver-dependent appearance properties, it is obvious that we can have nonperceptual thoughts about such properties. I can, for example, close my eyes and think about those properties, without thereby perceiving them. Thus the proposal cannot be simply that the contents of perceptions involve

---

example, Reid writes, "Here, then, is a phaenomenon of human nature, which comes to be resolved. Hardness of bodies is a thing that we conceive as distinctly, and believe as firmly, as any thing in nature. We have no way of coming to this conception and belief, but by means of a certain sensation of touch, to which hardness hath not the least similitude; nor can we, by any rules of reasoning, infer the one from the other. The question is, How come we by this conception and belief?" (1764/1997, pp. 57-58). Reid can be seen as holding that perception of hardness is a belief in the existence of hardness caused by a tactile sensation of it. This view thus has in common with Glüer's account, which holds that perception is simply belief. But Reid's view importantly differs insofar as it countenances nonconceptual qualitative states—and so does not count as a variant of the kind of representationalism I am discussing here. I will ultimately defend a view similar to Reid's in chapter 5.

[52] For a similar account of perceiver-dependent properties, see, e.g, Tye 2000, p. 95.

[53] It is likely that Glüer has something like Shoemaker's account in mind because she references him when introducing phenomenal properties (2009, p. 311, fn. 26).

those properties, but that they involve them in some specific way—for example, perception represents visual appearance properties as taking up parts of the visual field. But whether or not perception represents such perceiver-dependent properties, I'll argue that even this proposal fails.

If the differences in content must be differences in the ordinary viewer-independent properties represented, the most natural way to account for the difference between thoughts and perceptions is to hold that what distinguishes, for example, the belief that there is a dog from the visual perception that there is a dog is that the former exhibits the simple content *that there is a dog*, whereas the latter exhibits a content that involves the color and visual shape of the dog as taking up portions of one's visual field (cf. Dretske's proposal discussed above).

But at this point it is worth noting that this is a bizarre view. This account holds that the visual perception that there is a dog exhibits not only the content that there is a dog, but also content involving the dog's color taking up some part of one's visual field. But it certainly does not seem that way introspectively. Moreover, it is strange to think that all perceptions involve these complex intertwined contents. As Sellars puts it, "It is surely odd to suppose that such rich conceptual episodes as that expressed by 'here is a red book on a brown table' are invariably accompanied by such thin conceptual episodes as that expressed by 'here is a red rectangle standing out from a brown background'" (1967, p. 13).

A more fundamental problem with this view is that a belief arguably can exhibit this or any content without thereby being a perception (cf. Sellars 1975, pp. 306-307; Coates 2007, pp. 14-16). Though it may not be a regular occurrence, there is no reason to think that one could not form a belief that exhibits a complex content involving colors and shapes in one's visual field. And having that belief arguably need not result in one's perceiving those colors and shapes. It is open to the conceptualist representationalist to deny this, however. After all, it may be

impossible to have certain beliefs. For example, it does not seem like one could genuinely believe that circles are squares. Likewise, the conceptualist representationalist might argue that it is simply impossible to believe certain contents without thereby exhibiting a perception.

But perhaps the worst difficulty for this view is the most common objection raised to it. It is often observed that there are aspects of perception that are not belief-like, insofar as they are not open to revision in light of one's other beliefs. If I believe that it's raining, but then come to believe that the ground is not wet, that second belief may cause me to reject my first belief. But perception does not seem to operate this way.

The point is most obvious in the case of visual illusions. Consider, for example, the so-called waterfall illusion (e.g., Mather et al 1998). After temporarily fixating on the downward motion of a waterfall, people often subsequently experience so-called motion aftereffects wherein stationary objects appear to move upward. But, naturally, people nevertheless believe that these stationary objects are not moving. So our perceptions represent the upward motion of the stationary objects, despite the fact that we believe this not to be the case. Evans labeled this feature the belief-independence of perception (1982, p. 123). As many have argued, that there are aspects of our perceptions that are not susceptible to revision in light of other attitudes suggests that perceptions do not exhibit conceptual representational contents only (e.g., Evans 1982, p. 123; Crane 1992, pp. 149-154).

This worry has often been raised, and Byrne offers a reply to it.[54] Speaking of the illusion evoked by the Müller-Lyer figure, wherein two lines of equal length appear to be discrepant lengths because of differently placed arrows on their ends, Byrne writes,

[54] Glüer does as well. She proposes that since the content of the perception of the object's moving up involves perceiver-dependent appearance properties, whereas the content of one's belief that the object is not moving involves ordinary perceiver-independent properties, there is

> Inconsistent beliefs are perfectly common… it might be, as Armstrong once suggested, that this belief is 'held in check by a stronger belief '. If perception constitutively involves belief, then this neatly explains the commonly agreed fact that in the absence of any reason for thinking otherwise, one believes, in the usual unimpeded way, that the lines are unequal (2009b, pp. 450-451).

Byrne's idea is that when one perceives that the lines of the Müller-Lyer figure are unequal, one's belief that the lines are of equal length is "held in check" by one's perception. But this reply is unpersuasive. It is true that holding inconsistent beliefs is common, but holding outwardly contradictory beliefs is not. And if perception involves asserting some content, it appears that the state outwardly asserts that an object is moving and the object is not moving. What's more, the fact that it appears that the object is moving does not seem to be open to revision, even in principle. No amount of cajoling persuades the observer see the lines of the Müller-Lyer figure as the same length. What the conceptualist representationalist needs, then, is an explanation of why the contents of perceptual states, unlike other beliefs, so stubbornly resist revision. But there is no proffered explanation for this that does not seem like a thinly disguised effort to save the theory.

This criticism of the doxastic account of perception might be used to motivate the view that the perceptual attitude is a kind of perceptual attitude sui generis. The fact that perception involves a distinct attitude, one might argue, renders the perception immune to influence from other beliefs. But the criticism can be extended to apply to this view as well. Whatever the perceptual attitude is, it is assertoric. But when one asserts some content that content is typically open to revision in light of one's other attitudes. And, again, aspects of perception are not revisable in this way.

---

no conflict regarding their contents and there is no reason to suppose one's belief about the latter would influence one's perception of the former. Again, this proposal stands and falls on the controversial assumption that perception represents appearance properties.

Instead, one might think that an appeal to certain features of cognitive architecture might be of some help to the conceptualist here. Many have maintained that perceptual faculties are modular insofar as they are, for example, domain specific and informationally encapsulated (e.g., Fodor 1983). That is, the information contained within a module is proprietary to that module and not open to modification by other systems. By contrast, many hold that beliefs are inferentially integrated and flexible because they are part of a nonmodular central system, which receives information from, but cannot send information to, the perceptual modules. If the mind is functionally organized in this way, then one might argue that, even if perceptions exhibit many or even all of the same representational properties as beliefs, perceptions' location within the modules explains why beliefs have no purchase on them. But this proposal only pushes the question one step back. If one maintains that perceptions are modular, we need some explanation of this fact. And, for reasons that I have canvassed here, the conceptualist cannot appeal to either the states' contents or attitudes to explain why they play these special functional roles. Simply asserting that some beliefs (if the perceptual attitude is belief) or that some representational states (if the perceptual attitudes are sui generis) exhibit modularity does not explain the peculiar fact that such states are immune to revision in light of other beliefs.

I have argued that conceptualist representationalism cannot succeed. Again, I recognize that the argument just given fails to establish a version of representationalism as against the traditional qualia view. After all, the traditional qualia theorist has a neat explanation of the difference between perception and thought: Perception, but not thought, exhibits qualia. For now, however, it suffices to have demonstrated that there must be more to perception than conceptual representational properties.

§2.5 Atomistic Representationalism

Many representationalists do endorse the view that the contents of qualitative states are nonconceptual. But whether or not the contents of qualitative states are nonconceptual, all representationalists require an explanation of how such qualitative representation works. And as previously noted, representationalists, if they provide an account of the contents of qualitative states, typically endorse a version of a tracking theory of representation. On this view, for example, a perception or sensation exhibits a reddish qualitative character because it tracks red.

Any version of representationalism that depends on tracking semantics is atomistic insofar as a qualitative state's content—and thus its qualitative character—does not depend on that state's relation to any other mental states. Instead, the state's content depends only on its appropriate tracking relation to the perceptible property it represents there to be. It is thus consistent with atomistic representationalism that an individual not only perceive a single color at a time, but also be able to perceive only one color in its lifetime. Atomism about qualitative content echoes atomistic theories of conceptual content insofar as the latter hold that a thought's conceptual content depends solely on its causal connections to what it represents there to be, not on its relations to other thoughts.

There are various ways to understand the relevant notion of tracking as it figures in the tracking theory of content that undergirds atomistic representationalism. Brian Cutter and Tye offer the following relatively general characterization of it: "Tokens of a state S in an individual x represent that p in virtue of the fact that: under optimal conditions, x tokens S iff p, and because p" (Cutter and Tye 2011, p. 91; see also, e.g., Tye 2000, p. 136).[55] On this way of

[55] This characterization of the tracking theory entails that what states represent there to be are something like states of affairs or facts (states represent "that p"). One might object, however, that perception represents there to be objects or properties only (e.g., Rosenthal 2012, p. 24). I

understanding the tracking theory, for example, a qualitative state represents (that is, tracks) some perceptible property such as perceptible red if and only if, under optimal conditions, one tokens that state when perceptible red is present and because perceptible red is present.[56]

Considerable work is performed, of course, by the notion of optimal conditions. This component of the tracking theory is designed primarily to explain cases of misrepresentation. It is plain that one can represent there to be perceptible red even if perceptible red is not present, as in cases of hallucination, or fail to represent perceptible red even if red is the cause of one's perception, as in cases of ordinary red-green colorblindness. Tracking theories accommodate the possibility of misrepresentation by holding that it occurs when conditions are not optimal. For example, the reason why a person with ordinary red-green colorblindness does not veridically perceive perceptible red, even when perceptible red is present and the cause of one's perception, is that one is not in a state that has the biological function of tracking perceptible red.

Cutter and Tye gloss optimal conditions in the following way: "In the case of visual experience, we are in optimal conditions just in case we are in conditions of the sort that our visual system was designed to operate in by natural selection or by analogous processes in the course of ontogenic development" (2011, p. 91). Thus tracking theories of intentionality typically involve a teleological element.[57] On many versions of tracking theory, perceptions

---

will remain neutral for now about what kinds of things perceptions represent there to be. I will have more to say about this in chapter 5.

[56] As Cutter and Tye note, there are various versions of tracking theories that vary slightly in their details. For example, according to Dretske's so-called indicator semantics, a perception represents something if and only if the perception has the function to indicate that thing, wherein the notion of indication is understood in terms of representations whose functional roles are determined by natural selection (see, e.g., Dretske 1995, chapter 1). But the differences between these various views are not significant for present purposes (cf. Cutter and Tye 2011, p. 108, fn. 12).

[57] Not all tracking theories involve teleological elements. For example, according asymmetric dependence theory, a state represents something just in case certain counterfactuals are true of

represent what they do because they have the biological function, determined by natural selection or some relevantly similar process, to track those things.

Most representationalists assume an atomistic semantics for the view, I think, largely because holistic accounts of conceptual content such as CRS are often thought to face insuperable difficulties (see, e.g., Fodor and Lepore 1992). While I believe that holistic theories of conceptual content have much to recommend them, I will remain neutral about what theory explains the contents of conceptual states. But fundamental to any account of perceptual representation is an explanation of our perceptual discriminatory behaviors—and in chapter 5 I'll argue that the theory of qualitative content that best explains these and other phenomena is holistic. Thus my goal here is to argue that atomistic representational is problematic.

§2.5.1 A Commonsense Reason to Doubt Atomistic Representationalism

I'll begin with a commonsense reason to doubt atomistic representationalism. It's plain that we perceive colors to resemble and differ from one another in characteristic ways. For example, we perceive red to be more similar to purple than it is to green. Following Byrne and David Hilbert (2003, p. 13), I'll call the arrangement of resemblance relations among the colors the phenomenal structure of the colors. According to atomistic representationalism, what makes a perception exhibit the qualitative character that it does is its individual connection to the relevant perceptible property. Thus the theory provides no reason to expect that the colors would be perceived to relate to one another in characteristic ways.

Atomistic representationalists are, of course, well aware of this fact about perception. And they might reply that it poses no problem for their view. After all, they might argue, a

---

the state's relation to that thing. But the differences between these various views are not significant for present purposes.

theory of qualitative content need not explain everything about qualitative states—its only goal is to explain why a qualitative state represents a particular property, not necessarily why there are certain perceived similarities between properties. Moreover, they might urge, the phenomenal structure of the colors can be explained not by the theory of the semantics of qualitative content, but by the fact that qualitative states represent relational properties.[58]

But there are a couple of things to observe here. First, even if atomistic representationalists are correct that color vision represents relational color properties, this is an additional theory—a theory of what kinds of properties color perception represents there to be—that is needed to supplement atomistic representationalism's psychosemantics. Second, this additional theory is controversial. While it is phenomenologically obvious that red looks more like orange than it does like green, it is controversial that we perceptually represent colors to be relational properties. Even if colors are relational properties—a claim which is itself controversial—that fact alone does not entail that we must perceive them as such. It may be, for example, that time is a relational property, but it is not clear that we perceive it as relational. So the atomistic representationalist requires a reason to think that color perception represents colors as relational beyond the mere fact colors are relational.

My claim is not that atomistic representationalism cannot supplement its tracking semantics in some way to explain the phenomenal structure of the colors. Rather, my argument takes the form of an inference to the best explanation. In chapter 5, I'll offer a version of representationalism that depends on a theory of qualitative content that predicts and clearly explains the phenomenal structure of the colors on its own—and that does not advert to any

[58] This is essentially the strategy that Byrne and Hilbert themselves pursue; they argue that we do not represent individual colors per se, but rather what they call 'proportions of hue magnitudes' (2003, p. 14).

contentious theory of how we represent the colors in order to do so. On general theory-choice grounds, then, the theory that I'll propose is to be preferred to atomistic representationalism.[59]

Consider, by analogy, the debate between atomistic and holistic accounts of conceptual content such as CRS. Although everyone agrees that thoughts stand in various inferential connections to one another, atomistic theories neither predict nor explain this fact. By contrast, since holistic theories such as CRS explain conceptual content in terms of the inferential connections that hold between thoughts, such theories both predict and offer a tidy explanation of why such connections hold. Naturally, such considerations do not decisively decide between atomistic or holistic theories, but it is a point in favor of holistic accounts.

§2.5.2 Theoretical Reasons to Doubt Atomistic Representationalism

There are, in addition, theoretical reasons to doubt atomistic representationalism. First, as we recall from the debates over tracking semantics in the 1980's, it is notoriously difficult to spell out the appropriate tracking relation. Though there are many versions of tracking psychosemantics available, they all arguably face serious difficulties (for a review of various tracking theories and their problems, see, e.g., Neander 2006).

For example, since many versions of tracking psychosemantics depend on the teleological notion of biological function, such accounts are open to a well-known objection, which depends on Davidson's (1987) famous thought experiment involving the creature Davidson calls Swampman. Davidson imagines that a fortuitous lighting strike in a swamp could result in the development of a molecule-for-molecule duplicate of an ordinary person, the

[59] Indeed, the more atomistic representationalists must build into their accounts of how we represent the colors to explain the phenomenal structure of the colors, the more it may look as though the view is simply a notional variant of the holistic account that I'll sketch in chapter 5.

creature Swampman. It would seem, at least prima facie, that Swampman can enjoy thoughts and perceptions just like those of his normal counterpart. But Swampman's mental states do not exhibit biological functions as he has no evolutionary history—and so his states cannot meet the requirements of a teleosemantic theory of content. Thus, the objection goes, either tracking theories must hold that Swampman is in no representational states, which is highly questionable, or tracking theories are false (e.g., Neander 2006, pp. 384-384).

Unsurprisingly, atomistic representationalists have provided various replies to this Swampman objection and the literature is enormous. For example, Dretske (1995, pp. 144-149) has argued that our intuition that Swampman perceives anything is questionable—and so he simply denies that Swampman does perceive. By contrast, Tye (2000, p. 141) proposes that Swampman exhibits qualitative states, provided that the tracking theory of content is suitably adjusted to include additional counterfactual conditions. But such replies miss the point.[60] The bottom line is that our commonsense conception of qualitative states holds that we attribute them to a creature if it displays occurrent, functional, capacities to discriminate among perceptible properties. After all, the purpose of ascribing mental content to creatures in the first place is to explain and predict their actual behavior. If Swampman discriminates the colors in a sufficiently complex way that can be explained and predicted by ascribing perceptions to him, then common sense has it that he perceives the colors. A creature's history simply does not figure in our ordinary judgments about its capacity to perceive. We could follow Dretske and deny this central aspect of our everyday conception of qualitative states, but that is a big bullet to bite. At the least, representationalists should remain open to alternative theories of qualitative content—

[60] For reasons of space, I will not explore these or other atomistic replies to the Swampman objection (for discussion, see, e.g., Adams and Dietrich 2004).

and I'll develop shortly a theory of qualitative content that captures our ordinary conception of it.[61]

Recently, Pautz (e.g., 2006, 2010) has put forth another theoretical argument against atomistic representationalism, an argument that I find dispositive. Pautz's argument relies on considerations of what he calls hypothetical coincidental variation cases (2010, p. p. 337). Such cases involve two individuals who stand in the same optimal tracking relation to some percptible property, but, because of differences in their downstream neural processing of the information about that property, engage in radically different discriminatory or sorting behaviors. Consider, for example, Pautz's case involving Maxwell and Twin Maxwell (2006, p. 214). Maxwell is a normal human perceiver. Twin Maxwell lives on a planet much like Maxwell's, though the environment is slightly different.

Suppose that Maxwell and Twin Maxwell perceive an object that exhibits the same spectral-reflectance property, S, a property that we say is in the perceptible-orange range. And suppose, by stipulation, that both stand in the very same optimal tracking relation to S. Maxwell, of course, engages in various behaviors such as judging S to be more similar to properties that we hold are in the red range than they are to properties that we regard to be in the green range. Due to Twin Maxwell's alternative environment and thus alternative evolutionary history, he has developed a very different post-receptoral neural processing of stimuli than Maxwell's. When Twin Maxwell perceives S, he engages in very different discriminatory behaviors than Maxwell. Twin Maxwell, for example, judges the property to be more similar to properties that we hold are

[61] Atomistic theories of content that do not involve teleological components, such as asymmetric dependence theory, need not give the wrong result in the Swampman case—and so may seem to be preferable. But as I'll argue shortly, atomistic theories face other difficulties, difficulties which do not depend on their involving teleological considerations.

in the perceptible-green range than they are to properties that we regard to be in the perceptible-red range.

Pautz claims that a neutral description of the case suggests that, when both perceive S, Maxwell exhibits a perception that exhibits mental orange, whereas Twin Maxwell exhibits perception exhibits mental green. But, from the perspective of atomistic representationalism, Maxwell and Twin Maxwell must be undergoing qualitatively identical perceptions. Again, by stipulation, both Maxwell and Twin Maxwell both track S.[62] Thus atomistic representationalism is false. As the commonsense argument against atomistic representationalism above also revealed, something about how one sorts or discriminates a property from other perceptible properties is relevant to the qualitative character of one's experience. Mere tracking is not sufficient to determine qualitative character.

Atomistic representationalists have offered various rejoinders to this case (e.g., Byrne and Tye 2006, Lycan 2006). For example, the atomistic representationalist might maintain that Maxwell and Twin Maxwell do not both stand in optimal tracking relations to S, and thus do not exhibit qualitatively identical experiences (e.g., Byrne and Tye 2006, p. 252). One reason to think that, for example, Twin Maxwell does not stand in an optimal tracking relation to S is that he misrepresents S. After all, it would seem that Twin Maxwell perceives an object exhibiting S to be a shade of green, when in fact it is orange. So if Twin Maxwell is misrepresenting S, then he must not be tracking S in optimal conditions—thus atomistic representationalism is safe from Pautz's challenge.

[62] This is why the case is not akin to an ordinary case of color blindness. If Twin Maxwell were simply colorblind, then there would be no problem for atomistic representationalism. But Twin Maxwell is stipulated to be a normal perceiver whose visual system has the biological function to track S and conditions are optimal.

But, as Pautz emphasizes, Twin Maxwell is stipulated to be tracking S in optimal conditions. Moreover, the fact that Twin Maxwell is misrepresenting S is not grounds to deny that he is tracking it in optimal conditions. Tracking theorists such as Tye characterize veridical representation independently in terms of optimal conditions, so it would be question begging to claim that optimal conditions fail to obtain whenever misrepresentation takes place (cf. Pautz forthcoming, p. xxx). Without an independent reason to think that conditions are not optimal, atomistic representationalism delivers the wrong result that both Maxwell and Twin Maxwell represent there to be the same perceptible property and so exhibit the same qualitative characters. Thus, again, representationalists should be open to an alternative theory of qualitative content.

§2.5.3 An Experimental Reason to Doubt Atomistic Representationalism

I'll close this section by mentioning an experimental reason to doubt atomistic representationalism. Consider the well-known evidence of so-called color-context effects—that is, the evidence that what colors are in the environment of a particular color affect the way that color is perceived (for an artful discussion of cases with many examples, see, e.g., Albers 2006). In the familiar simultaneous-brightness contrast illusion, for example, a shape of a given shade of blue is illusorily perceived to be a brighter shade of blue when placed against a dark background than the very same shade of blue is perceived to be when placed against a lighter background (see, e.g., Hardin 1993, p. 24).[63] Although such color-context effects are commonplace, we rarely appreciate them—and noticing them can be quite striking. Within some bounds,

[63] It is worth noting that this effect is not only quite robust for conscious perception, but has recently been observed to occur in cases of nonconscious perception (see, e.g., Persuh and Ro 2011).

practically any color can be perceived as some other color, provided its background colors are suitably adjusted.

As I argued in connection with the facts regarding the phenomenal structure of the colors, atomistic representationalism does not predict these experimental findings nor does it have the resources on its own to explain them. Again, according to atomistic representationalism, a state's qualitative character is determined by its individual connection to the relevant perceptible property. Thus the theory provides no reason to expect that the way a perceptible color appears would be susceptible to influence by the perceptible colors in its surround.

As was the case with the commonsense evidence presented above, atomistic representationalists are well aware of this kind of experimental evidence. And once more, they might reply that a theory of qualitative content need not explain why there are these kinds of effects. Instead, they might argue that color-context effects are merely a product of the processing of our visual systems (e.g., Tye 2000, p. 156). But any theory of qualitative content must grant that the way colors appear can be explained by the way our visual systems operate. The point is that atomistic representationalism itself sheds no light on the phenomena of color-context effects. Again, my argument here is abductive. In chapter 5, I'll offer a theory of qualitative content that suggests why such effects might occur and provides the beginnings of an explanation of them—and so is again to be preferred to atomistic representationalism on general theory-choice grounds.

§2.6 Conclusions

Most qualophiles object to representationalism on the grounds that it renders impossible undetectable quality inversion without misrepresentation. As Shoemaker (2003, p. 255)

observes, because standard representationalists hold that there can be no change in a qualitative state's mental properties without a change in the perceptible properties that the state represents there to be, they hold that undetectable quality inversion without misrepresentation is impossible (see, e.g., Tye 2000, p. 110, Speaks 2011).[64] Tye, for example, is happy to grant that the scenario is conceivable, but he denies it is possible. Tye writes,

> So long as these cases are conceptually possible, whatever their metaphysical status, the pure representationalist is on secure ground in allowing the possibility of rife intraspecies inverted spectra without genetic defects…. Just as we can coherently think that Hesperus is not Phosphorus or that water is not H_2O, so, too, we can coherently think that sensory misrepresentation is rife among biologically normal humans, even if these things are all metaphysically impossible (2000, p. 110).

Recently, however, several representationalists have proposed versions of the view that purport to accommodate this possibility. In the following chapter, I will explore three of these views in detail.

[64] For example, Harman writes, "I am not sure that the imagined possibility of inverted spectra is really coherent. (Please note: I am not sure.) A red object supposedly looks Q to George and T to Mary. It would seem that an object cannot be both Q and T in the same place at the same time in the same way. That would be like an object's being both red and green at the same place and time in the same way. But then, either George's experience or Mary's experience or both of their experiences must be in some respect nonveridical, incorrectly representing the object seen" (1996a, p. 12).

CHAPTER 3: REPRESENTATIONALISM AND QUALITY INVERSION

§3.1 Introduction

As noted in chapter 1, the traditional qualia view—the view that mental qualities are nonrepresentational and perhaps resist informative explanation altogether—is largely motivated by various intuitions regarding mental qualities. One prominent intuition of this sort is the intuition that undetectable quality inversion without misrepresentation—what I'll call 'quality inversion' for short—is possible or at least conceivable (e.g., Sextus Empiricus 2005, Locke 1700/1975, Shoemaker 1982, Block 1990). According to this intuition, it is possible that two people can veridically perceive the same object and yet exhibit distinct mental qualities, but in a way that cannot be detected even in principle. Many theorists in the contemporary philosophy of mind believe that the starting point of any account of mental qualities is to accommodate this sort of intuition. If this is the appropriate route to approach mental qualities, then it may be that they cannot be explained in any naturalistic or informative way. Perhaps there really is a hard problem.

Nevertheless, some representationalists maintain not only that we can accommodate quality inversion, but also that we can give at least a partial account of mental qualities in terms of their roles in representation. In this chapter, I'll consider three such proposals: Shoemaker's (1994b, 2000, 2001, 2006) dispositionalist representationalism, Chalmers's (2004, 2006, 2010) so-called Fregean representationalism, and Block's (1996, 2003, 2007a, 2007c, 2010) account of mental paint (a version of weak representationalism). According to the holistic version of representationalism that I'll develop in chapter 5, by contrast, quality inversion is impossible.

Since the views that I'll discuss in the present chapter accept the possibility of inversion, the complex machinery that these proposals employ is fabricated to support a questionable datum

at the outset, and is in that way unmotivated. I'll argue, moreover, that there are independent reasons to think that none of these accounts can explain how quality inversion works. These accounts thus fail to put forward credible alternatives to the version of representationalism that I'll develop. Indeed, I'll argue that these accounts fare no better than the traditional qualia view's concession that mental qualities are completely mysterious.

§3.2 Shoemaker's Dispositionalist Representationalism

In the philosophy of perception, perhaps the two most commonly held intuitions about qualitative character are often thought to be contradictory. On one hand, there is the intuition that quality inversion is possible. Other the other hand, many have had the intuition that perception is, as G. E. Moore (1903) put it, transparent or diaphanous (p. 446 and p. 450 respectively; for modern statements, see, e.g., Harman 1990; Tye 1995; Jackson 2007). Though there is debate about the nature of this latter intuition (see, e.g., Nida-Rümelin 2007), it roughly holds that in introspection we see right through our perceptions, as it were, to the perceptible properties of objects that we perceive there to be. If one examines one's perception of a ripe tomato, one is invariably aware of the redness of the tomato itself.

These intuitions are typically thought to be incompatible for the following reason. The transparency intuition is often taken to entail a form of representationalism. If our awareness of our states' qualitative characters always involves an awareness of the qualities that we represent there to be, then it would seem that representationalism follows.[65] But standard versions of

[65] I will not examine this argument for representationalism in detail in this dissertation because I think it is, at bottom, question begging. As several proponents of the traditional qualia view have observed, the transparency intuition is a phenomenological observation—and it is open simply to deny it. Thus many qualophiles claim that they are often, if not always, aware of

representationalism are incompatible with the possibility of quality inversion with (see, e.g., Harman 1996, p. 12; Tye 2000, p. 110; Speaks 2011). If there can be no difference in two people's perceptions without a difference in the properties those states represent there to be, then two people cannot veridically represent there to be the same things and yet exhibit distinct qualitative characters.

No one has made a more valiant attempt to reconcile these intuitions than Shoemaker. In a series of papers, Shoemaker (1994b, 2000a, 2000b, 2006) has developed and defended one of the most innovative and complex accounts of perception to date. Shoemaker describes the following as his paradigm case of quality inversion:

> A case of spectrum inversion, of the sort that we are interested in, must be one in which the invertees do not differ in their beliefs about the objective colors of things [and do not misrepresent those objective colors]…. If Jack and Jill are spectrum inverted relative to each other, this will involve, for example, the way red things look to Jack being the same as the way green things look to Jill and vice versa… (2006, pp. 463-464).

According to this way of understanding inversion, Jack and Jill can both perceive a ripe tomato to be red, but the tomato looks red to Jack and looks green to Jill. That is, the tomato looks to Jill the way green limes look to Jack. But what is it for an object to look red in addition to its being perceptible red? It is doubtless that 'looks', as well and related expressions such as 'appears' and 'seems', is ambiguous in several ways (see, e.g., Chisholm 1957, chapter 4; Austin 1962, chapter 4; Jackson 1977, chapter 2).[66] But there is arguably a sense in which an object's looking some way to someone amounts to that person's representing that object to be some way. Shoemaker (1982, p. 365) has called this the intentional sense of that kind of expression. On this

---

nonrepresentational qualia in introspection (see, e.g., Block 2003, Burge 2003, Kind 2003, Prinz 2012).

[66] For example, Chisholm distinguishes what he calls the comparative use of appearance words—as in, Jack appears taller than Jill—from what he calls the epistemic use—as in, it appears as though Jack is out of town (from the fact that his house is empty).

reading of 'seems', an object's seeming some way to someone amounts to that person's representing that object to be that way (see, e.g., Schellenberg 2011, pp. 722-723 for a defense of this reading).

But if Jack and Jill can both veridically represent the tomato to be perceptible red, then they must also be representing it as exhibiting properties distinct from its perceptible color. These distinct properties, which Shoemaker calls appearance properties, are also perceived in addition to an object's perceptible properties.[67] However, unlike perceptible properties, appearance properties are mind-dependent properties that cause experiences with certain mental qualities. As Shoemaker puts it,

> To a first approximation, an object's having [an appearance] color property just is its looking a certain way to certain perceivers in virtue of having a certain color, and this normally amounts to the color of the object presenting itself in one of the ways it can present itself (2000a, p. 466).

Thus in addition to denoting perceptible colors and mental qualities, predicates such as 'red' also denote appearance properties. I will signal when such predicates denote appearance properties by affixing 'appearance' to them. So a tomato looks to be objective red to Jack because he also perceptually represents it as appearance red, and thereby his perception exhibits a red quale. By contrast, the tomato looks to be objective green to Jill because she represents it as exhibiting appearance green and thereby her experience exhibits a green quale.

Shoemaker maintains that we do not enjoy direct introspective access to our qualia. In that way, he accommodates the transparency intuition. But Shoemaker claims that we nonetheless know that such mental qualitative properties exist, based upon our reflections on the

[67] Shoemaker initially calls these properties phenomenal properties in his (1994b) and (2000), but calls them appearance properties in subsequent publications. I will refer to them throughout as appearance properties. For other accounts of appearance properties, see, e.g., Tye (2000, p. 95), Noë (2004, p. 83), Glüer (2009, p. 298), Hill (2009, p. 139).

conceivability of inversion. On Shoemaker's view, it is in virtue of being directly aware of the appearance properties of objects that we can be inferentially aware that mental qualities exist and that they can be inverted.[68]

At first blush, it may seem as though Shoemaker's ontology of perceptible colors, mental colors, and appearance colors posits too much. Similarly, it strikes many as odd that one can exhibit mental qualities, but never be introspectively aware of them.[69] But for Shoemaker, these various kinds of properties are necessary to explain our ordinary intuitions about perception. And mental qualities are theoretical entities posited to explain appearance properties in a noncircular way. Shoemaker writes,

> The relation of qualia to this phenomenal character is not that of being it, and not that of having it, but rather that of being constitutive determiners of it…. It is partly in virtue of having the qualia it does that the experience represents what it does…. And part of what it represents is the instantiation of a property, a "phenomenal property," which is in fact, although it is not explicitly represented as, the relational property of producing, or being apt to produce, experiences having these mental qualities (1994b, p. 29).[70]

Shoemaker thus develops a version of representationalism, which is putatively compatible with the possibility of quality inversion. On Shoemaker's view, there can be no change in a state's mental qualities without a change in its representational content. But Shoemaker's view is not a

---

[68] As Block puts it, "according to Shoemaker, we have a kind of indirect introspective access to the appearance character of the experience via our intuitions about the inverted spectrum" (2007, p. 541). Though Shoemaker does not use these terms, another way to express this notion is that we have a sort of displaced perceptual access to our qualia via our reflection on the possibility of quality inversion (see, e.g., Dretske 1995, p. 41).

[69] For example, using 'mental paint' to refer to mental qualities, Block argues, "Shoemaker's view is highly paradoxical—there is mental paint but it is a theoretical entity. It is of the essence of mental paint to be something of which we are aware" (2007, p. 541). Block is correct that Shoemaker's view does not capture this commonly assumed feature of inversion. In any case, I'll argue that Shoemaker's view fails to accommodate any kind of inversion.

[70] On Shoemaker's view, mental qualities are nonrepresentational properties of mental states, which can be identified and individuated by their relative locations in spaces of mental qualities. But these details matter at this point. I will discuss Shoemaker's approach to mental qualities in more detail in chapter 5.

version of standard representationalism, insofar as he denies that the qualitative character of a state simply is (or is determined by) the representation of ordinary perceptible properties. On Shoemaker's view, two qualitatively distinct perceptions can veridically represent the same perceptible color because those states also represent distinct appearance properties responsible for causing those different qualia.

Shoemaker's view is often described as a dispositionalist version of representationalism insofar as appearance properties are often understood to be dispositional properties. But over the years, Shoemaker has refined his characterization of appearance properties in response to various objections leveled against it. Recently, several have argued that, regardless of how Shoemaker understands appearance properties, his account cannot accommodate spectrum inversion (e.g., Chalmers 2004; Thompson 2007). And some have offered Shoemaker-inspired accounts to attempt to correct these problems (e.g., Egan 2006).

In this chapter, I'll argue that the problem with Shoemaker's view is not his account of mental qualities per se, but its treatment of perceptible properties. It is a folk-psychological truism that, normally, perceptible red things look red: More accurately, if something is, for example, perceptible red, then it looks under standard lighting conditions to normal observers to be perceptible red (see, e.g., Sellars 1956/1997, p. 48). Call this ordinary claim about perceptible colors the perceptible-color truism.[71]

The perceptible-color truism does not entail that the perceptible colors need be somehow defined in terms of the way things look or that perceptible colors are somehow mind dependent. Rather, the truism simply reflects the fact that, to use Saul Kripke's expression (e.g., 1980, p.

[71] The following discussion focuses on perceptible color, but the same considerations arguably apply to the perceptible properties sensed by other sensory modalities such as smells, tastes, and sounds.

55), we fix the reference of our perceptible color terms by reference to the ways perceptible red things look under standard lighting conditions to normal observers.[72] Even if we fix the referents of our terms this way, the perceptible colors can nonetheless be perfectly respectable mind-independent physical properties.

But Shoemaker rejects the perceptible-color truism. According to Shoemaker, an object can, under standard lighting conditions, look very different to normal observers. A ripe tomato is perceptible red and yet, while it looks red to Jack in broad daylight, it simultaneously looks green to Jill. Since Shoemaker abandons this commonsense way of fixing the referents of our terms for perceptible colors, he owes an account of the nature of perceptible properties. However, the basic problem for Shoemaker, I'll argue, is that he cannot provide an informative account of perceptible properties such as perceptible colors.

One might immediately reply that the fact that Shoemaker cannot provide an informative account of the perceptible colors is not a problem for him. After all, isn't eliminativism about color a possibility? But Shoemaker is an avowed realist about colors; he writes, "I assume… that objects in the world have colors, and have them independently of being perceived to have them, and independently even of there being creatures capable of perceiving them" (2003, p. 253). And Shoemaker must be a realist about perceptible properties. If Shoemaker were to eliminate the perceptible colors, he would not be able to accommodate one of his central desiderata—namely, undetectable spectrum inversion without misrepresentation. If there are no perceptible colors, then we all systematically misrepresent the world as being colored.

In addition, Shoemaker claims that, in perceiving an object such as a tomato, one perceives not only its appearance color but also its perceptible color. In particular, Shoemaker

[72] Issues can be raised around how to identify standard lighting conditions and normal observers, but these questions are not relevant here.

claims that we see an object's color in virtue of seeing its appearance color. Speaking of appearance and perceptible properties, Shoemaker writes,

> [T]hese are not independent and separable aspects of the experience's content; rather, the experience represents the color by representing the appearance property. To put it otherwise, we see the color of a thing by seeing an appearance property it presents (1994b, p. 35).

For these reasons, Shoemaker owes an account of the perceptible colors, or at least must make room for such an account. That is, Shoemaker must provide (or leave open the possibility of) an account that enables us to explain how we see an object's color in virtue of seeing its appearance color. But I'll argue that Shoemaker's view cannot accommodate an account of the perceptible colors and is thus unworkable. My argument here will be distinct from Tye's well-known objection that Shoemaker renders it impossible to see the perceptible colors because they are placed behind a veil of appearance colors (see, e.g., Tye 2000b, p. 464; for a reply, Shoemaker 2000a, p. 466). My claim is not that we cannot perceive an object's perceptible color in virtue of perceiving its appearance color. Rather, I'll argue that, by distinguishing these kinds of properties, we have no explanation of how this could be the case.

Shoemaker himself does not present a theory of the perceptible colors, but there does seem to be a promising route available. Michael Thau (2002) develops a theory of appearance properties much like Shoemaker's—and Thau suggests that each perceptible color can be understood in terms of their relations to appearance colors. Thau proposes that "each color C can be defined as a disposition to cause a representation involving the property perceptually correlated with C" (2002, p. 236). In other words, each perceptible color will be correlated with a unique set of appearance colors and that each perceptible color can be identified with the disposition to cause that unique set of appearance colors. If so, Thau argues (2002, p. 236), then we not only have a theory of the perceptible colors, but we also have a straightforward

explanation of how one represents an object's perceptible color in virtue of seeing its appearance color. However, I'll argue that on any of Shoemaker's and others' successive accounts of the appearance colors, there is no way to provide an account of the perceptible colors in terms of the appearance colors.

§3.2.1 Shoemakerian Iterations

Before discussing the various ways in which Shoemaker and others have glossed appearance colors, I should mention that I do not examine in great detail the motivations underlying each of the successive accounts. My point is that, regardless of how we understand appearance properties, there can be no account of perceptible properties.

In one of his early attempts at characterizing appearance properties, Shoemaker claims that an appearance property is "the property of producing an experience with a certain qualia in something related in such and such a way to the possessor of the property" (1994b, p. 35). According to this proposal, for example, appearance red is the property of currently causing mental red in at least some perceiver under some lighting conditions. But, as Shoemaker himself later observed, this proposal entails that an object has no appearance colors when unseen. Per his color realism, however, Shoemaker holds that objects have perceptible colors even when unseen. The first version of the account therefore entails there is no relationship between an object's color and its appearance colors. Objects can exhibit distinct perceptible colors even if they do not exhibit any appearance colors. If the only way to identify the colors is in terms of distinct sets of appearance colors, this account is unworkable.

Shoemaker (2000) himself found troubling that objects do not exhibit their appearance colors when unseen and proceeded to revise his proposal to conceive of appearance properties as

dispositional properties. After all, an object can exhibit a dispositional property without that disposition ever being triggered—a fragile vase may never break. Thus, if appearance colors are dispositional, then objects can exhibit them even if no one has ever seen them to experience resulting mental qualities. But because Shoemaker recognized that such dispositions would not cause the same sorts of experiences in all perceivers, he characterizes the dispositions as relative to perceivers of certain sorts. Shoemaker writes,

> Each appearance property could be defined, not as a disposition to produce a certain sort of experience in all visual perceivers, and not as a disposition to produce such experiences in visual perceivers having a visual system with a certain makeup, but as a disposition to produce such experiences in creatures with visual systems of one or more sorts (2000, p. 467).

According to this revised view, appearance red is the disposition to cause mental red in perceivers of a given type under some lighting conditions. As noted before, Shoemaker never gives an account of the perceptible colors, but a natural suggestion would be to identify the perceptible colors with the categorical bases of these dispositions. Perhaps perceptible red, then, is the spectral-reflectance property that is the categorical base of appearance red. Identifying the perceptible colors in this fashion neatly explains how one perceives an object's perceptible color by perceiving its appearance color. Just as one can perceive a dispositional property, such as the fragility of a vase, so too through this perception one also perceives some categorical property of the thing that underlies the dispositional property.

In developing his Shoemaker-inspired account, Uriah Kriegel (2009, p. 96) takes a path along these lines, identifying the perceptible colors with the categorical bases of the dispositional appearance properties that he posits—namely, the dispositions to cause the realizers of mental qualities. Imagining a person he calls Norma, a human who has an experience exhibiting mental blue16, Kriegel writes that the appearance property that she perceives should be understood thus:

> The disposition to produce [neural] response N in Norma and subjects physically like her within her gender, race, and age group and to produce [silicon] response S in Venutian Norma and subjects physically like her in gender, race and age group and… and so on, for every type of realizer [of the experience of blue16]" (2009, p. 90).

The problem with this view, however, stems from the relationship between perceptible colors and appearance colors. Because Shoemaker countenances undetectable quality inversion, an object with almost any spectral-reflectance profile has the disposition to cause any mental color in perceivers of some type under some circumstances. If Jack and Jill are inverted twins who have different types of visual systems, then a surface that reflects 650 nm light (which we regard as perceptible red) looks to Jack the way a surface that reflects 500 nm (which we regard as perceptible green) looks to Jill, and vice versa. As a result, both the 650 nm and 500 nm surfaces exhibit the same set of appearance properties. Moreover, if we imagine creatures with visual systems different than ours, then it is easy to imagine that both the 650 nm and 500 nm surfaces could look any way at all, at least to some type of creature (perhaps within some bounds given the way that light combines and restrictions on how visual systems operate). So this view entails that every spectral-reflectance profile exhibits every appearance color. Thus there is no relationship between the perceptible colors and the appearance colors and there is no way to identify the former in terms of the latter.

Moreover, accounts which complicate things in this way result in another sort of problem. On this view, something's appearing blue amounts to my perceptually representing it as being the cause of my experience, but also as its being the cause of all sorts of other creatures' experiences. But as Pautz (2010a, p. 352) asks, what sort of psychosemantics could deliver the result that one perceives properties that make reference to alien perceptual systems? Kriegel himself endorses a relatively simple Dretskean teleological tracking theory of representational content, according to which a representation represents what it does if it has evolved the

biological function to track that thing (see, e.g., Dretske 1995). But how could it be evolutionarily advantageous that I perceptually represent anything about Venutians, with whom no one in my species has ever come (or will ever be likely to come) into contact? Indeed, it is implausible that any theory of intentionality could explain why the contents of ordinary perceptions of perceptible colors make reference to aliens.

In his (2001), Shoemaker tweaked his proposal to attempt to address these issues. Instead of claiming that appearance properties are first-order dispositional properties, he instead argues that they are higher-order dispositional properties. On this view, appearance red is not itself a disposition relative to perceivers or conditions. Rather, appearance red is the property shared by every object that is disposed to cause mental red in some perceiver under some conditions. Perhaps, along these lines, an object's perceptible color is the conjunctive property of exhibiting some set of appearance colors. What Shoemaker needs, then, is for each spectral-reflectance profile to correlate with a unique set of appearance properties—not construed as dispositions, but in this higher-order fashion. But a related problem for this view arises here as well. An object with any spectral-reflectance profile will exhibit every higher-order property of causing mental qualities of some sort in particular groups of perceivers under some conditions. Thus, once more, there is no unique relationship between a spectral-reflectance profile and any set of appearance properties.

In response to several other critiques, Shoemaker offered his most recent (2006) proposal, which holds that perception represents objects not as having appearance properties, but as what he calls qualitative characters. Shoemaker writes,

> A color will have different qualitative characters corresponding to the different ways that it can look (without being misperceived) in different viewing conditions, and also different qualitative characters corresponding to the different ways it can look to different observers owing to differences in their perceptual systems (2006, p. 474).

On this view, qualitative characters are higher-order properties of perceptible properties, properties which correspond to the all of the ways that a perceptible property can appear to a particular perceiver. Qualitative characters, unlike appearance colors, are not individuated by perceivers or lighting conditions. But the same fundamental problem plagues this account. Since any object can look any way to some perceiver under some conditions, this proposal too has the consequence that every object will exhibit every qualitative character. Even here, there is no way to individuate the perceptible properties in terms of qualitative characters.

None of Shoemaker's own iterations, I argue, make room for an informative account of the perceptible colors. But what of some Shoemaker-inspired views? After arguing that Shoemaker's own accounts cannot accommodate undetectable quality inversion, Andy Egan (2006) suggests that appearance properties are not, strictly speaking, properties. Rather, they are a kind of property-analogues, which he calls centering features (2006, p. 511). Egan cites the property of being nearby as an example to explain the nature of centering features. Roughly, if one represents something as being nearby, one arguably does not attribute the property of being nearby to that thing. The particular thing can be nearby to some things, but not nearby to others. And many things can be nearby at once. Thus representing something as nearby does not involve attributing a property to it, but involves representing it as having a location relative to your own location.[73]

Likewise, Egan's idea is that attributing appearance colors to things is to attribute these sorts of centering features. Egan writes,

[73] Egan unpacks this idea in terms of a complex view about possible-worlds contents and centered worlds (see, e.g., 2006, pp. 509-512), but examining this view would take us far afield from our discussion.

> Here is a first pass at a proposal for what [appearance green] is: it's the centering feature, being disposed to cause [mental green] experiences in me (or possibly, being disposed to cause [mental green] experiences in perceivers of my kind) (2006, p. 513).

Roughly, one's representing something as having a particular appearance color is to represent it as being the kind of thing that is disposed to cause a particular mental quality in one. Like the property of being nearby, this centering feature can be attributing to countless items, which may be different for different individuals at different times. So, for example, if I perceive an object that reflects 500 nm light, and that object causes in me mental red, then I perceive that object as exhibiting the centering feature appearance red. By contrast, if my inverted twin perceives that object and it causes in her mental green, then she perceives it as appearance green.

While this account may be of help when accounting for undetectable inversion, it is clearly no help with the problem I am raising. Since any object can look any way to some perceiver under some conditions, this proposal too has the consequence that every object will exhibit every appearance property, and thus there is no way to individuate the perceptible properties in terms of them.

Though he does not provide a distinctive account of appearance properties, Thau (2002) does defend a view much like Shoemaker's. Interestingly, Thau argues that we do not perceptually represent objective colors at all. Rather, on Thau's view, we perceive only appearance properties. Shoemaker has been a vocal critic of Thau's position (for reasons that I will not discuss[74]), although one might think that the above considerations would push Shoemaker to adopt something like Thau's view. However, this type of retreat is not helpful. Thau does not hold that there are no objective colors, only that we do not perceive them in virtue of perceiving the appearance ones. Indeed, Thau is clear that we do represent objective colors,

[74] See, e.g., Shoemaker (2006) and (2007).

but in downstream judgments and not in perception (e.g., 2007, p. 638). So even if Thau were right that we do not perceive objective colors, we would still require an account of the objective colors—and it is not obvious what sort of account there could be given his Shoemakerian approach.

§3.2.2 Shoemaker and Standard Theories of Color

I have argued that the most promising route to explain the perceptible colors—explaining them in terms of the appearance colors—is a dead end. But there are, of course, several theories of perceptible colors on the market—and perhaps one of these theories is available to Shoemaker. I have already argued that eliminativism about perceptible color is not a live option for Shoemaker. But what about a standard version of color physicalism, according to which perceptible colors are simply identified with underlying physical properties of surfaces such as spectral-reflectance properties (e.g., Reid 1764/1997; Tye 2000)? On this view, for example, we can identify the objective reds, as we already do, with the properties of reflecting wavelengths of light between roughly 635 nm and 700 nm. And the property of reflecting each individual wavelength, such as 635 nm or 636 nm, is identical to a distinct shade of perceptible red. On this view, a perceptible color is completely independent of any appearance colors that may be related to it.

But in conjunction with Shoemaker's approach this view of perceptible color yields bizarre results. Take some spectral-reflectance property within what we think of as the range of perceptible red, such as the property of reflecting light that has a wavelength of 650 nm. On most of Shoemaker's views, objects that exhibit this spectral-reflectance property exhibit both appearance red and appearance green. Now consider some other spectral-reflectance property

within what we think of as the range of perceptible green, such as the property of reflecting light that has a wavelength of 500 nm. Shoemaker holds that the perceptible color that reflects 500 nm light also exhibits both appearance red and appearance green. But which spectral-reflectance property, the property of reflecting 650 nm light or the property of reflecting 500 nm light, is perceptible red and which is perceptible green? By hypothesis, these are distinct perceptible colors. Bu if we are classifying perceptible colors in terms of spectral-reflectance properties, the answer would seem to be that the object that reflects 500 nm light is perceptible red. After all, perceptible red is perceived to be more similar to perceptible violet than it is to perceptible green, and the perceptible violet range is closer to 500 nm than it is to 650 nm. But that conclusion is wrong.

It is important to note that this problem is quite different than the oft-observed problem for standard physicalist theories of perceptible color, that there appears to be no interesting way in which the structure of the perceived similarities and differences between the perceptible colors matches the physical similarities that characterize those physical colors. This problem is often illustrated by the phenomena of so-called metamers: Members of a metameric pair of spectral-reflectance properties are perceptually indistinguishable under particular illumination conditions, and yet are different physical properties. On this basis of this and other evidence, many have concluded that color realism is false (see, e.g., Hardin 1993, pp. 63-64).

In reply, several color physicalists have suggested that we identify perceptible colors not with particular spectral-reflectance properties, but with relevant equivalence classes of them. As is well known, the human color-vision system operates according to an opponent-processing model, wherein three types of color-sensitive cones in the human eye are activated at different levels to produce visual experiences. According to this theory of the human visual system,

metamers are perceptually indistinguishable because they activate equivalent ratios of these color sensitive cells. It is therefore open, some argue, to identify the perceptible colors with equivalence classes of spectral reflectances that reflect distributions of wavelengths resulting in particular ratios of activation among the three color-sensitive cones in the eye under standard lighting conditions (e.g., Byrne and Hilbert 2003, p. 12; Rosenthal 2005, p. 141).

But even if we identify the relevant physical properties not with individual spectral-reflectance properties but with these sorts of equivalence classes of them, Shoemaker must allow that each equivalence class will nonetheless exhibit the same set of appearance properties. And thus he cannot provide an explanation of which individual spectral-reflectance properties (or classes of spectral-reflectance properties) correspond to which perceptible colors.

Perhaps there is an obvious option available to Shoemaker. He might instead endorse versions of color dispositionalism, according to which an object's perceptible color just is the disposition to look a certain a way to certain perceivers or groups of perceivers under certain circumstances (e.g., Johnston 1992; Levin 2000). Color dispositionalism is widely recognized to be a form of subjectivism about color because it builds observer-relativity into the nature of colors. This view of color might seem be in conflict with Shoemaker's avowed color realism. But all Shoemaker requires for his account of perception is an account of perceptible color that enables him to explain how we veridically represent perceptible colors in virtue of seeing appearance colors, whether or not those perceptible colors are subjective or perceptible in nature.

Indeed, color dispositionalism echoes Shoemaker's account of appearance colors in many ways and seems to fit with it nicely. According to a certain version of color dispositionalism, perceptible red is the property of looking red to some particular individual in some particular circumstance, to another individual under different circumstances, and so on. And according to

Shoemaker, a perceptible color's looking some way is explained in terms of its exhibiting an occurrent appearance property. The combination of a Shomakerian view of appearance properties and color dispositionalism about perceptible color, then, holds that perceptible red is the disposition to manifest occurrent appearance red to some particular individual in some particular circumstance, to another individual under different circumstances, and so on.

Interestingly, this view would seem to provide a way to individuate the perceptible colors. According to this view, though an object that is perceptible red and an object that is perceptible green will both exhibit appearance red and appearance green, only the object that is perceptible red will cause occurrent appearance red in Jack under normal lighting conditions, whereas only the object that is perceptible green will cause occurrent appearance red in Jill under normal lighting conditions. By relativizing perceptible color in this way, this view entails that each perceptible color can be individuated in terms of the unique set of occurrent appearance properties that it causes to manifest to particular observers under various conditions.

This account is intriguing, but there are various problems with it. First, the view does not provide any principled grounds for identifying a particular perceptible color with any given unique set of dispositional properties. The views do not explain, why, for instance, the spectral-reflectance property that is the categorical base of the disposition to cause mental red in Jack under normal lighting conditions is perceptible red. After all, the spectral-reflectance property referred to as 'perceptible green' in the previous paragraph also causes mental red in some normal observers under some normal conditions—viz., Jill. We could, of course, simply stipulate that perceptible red be identified with one physical categorical base of one unique set of dispositional properties, but the choice would be arbitrary.

Moreover, even if there were some non-arbitrary way to associate each unique disjunction of dispositional properties with a perceptible color, it is unclear how, on such views, one perceives an object's perceptible color in virtue of perceiving its appearance color. Suppose that perceptible red is the categorical base of the disposition to cause mental red in Jack under normal lighting conditions. And suppose also that Jack sees something perceptible red under normal lighting conditions and thereby exhibits mental red. How does Jack know that he sees something perceptible red? All Jack knows is that sees something that exhibits the disposition to cause mental red in him in normal lighting conditions. In order for Jack to know that the categorical base of this deposition is perceptible red, he must know that he is not a creature like Jill (that is, that he is not an invert), which he cannot know.

In view of the failures of these sorts of accounts, one might think that Shoemaker could instead take perceptible colors to be primitive and thus free from burden of explanation (cf. Campbell 1993/1997; Watkins 2005). But primitivism is not only mysterious and explanatorily unsatisfying, but problematic in various additional ways (see, e.g., Byrne and Hilbert 2007). And it is not obvious that Shoemaker can explain how we perceive perceptible colors by perceiving appearance colors, a process which demands its own explanation even if we take the former to be primitive. Indeed, a view that can explain mental qualities without eliminating perceptible properties or rendering them inexplicable is clearly to be preferred—and I'll develop such a view in chapter 5.

In sum, no Shoemaker-style account can accommodate quality inversion and provide an informative theory of the perceptible colors. If this central inspiration for the theory cannot be accommodated, then the rest of the theory is unmotivated. In chapter 5 I'll develop a theory of

mental qualities that both adequately accounts for them, while also making room for an informative account of perceptible properties.

§3.3 Chalmers's Fregean Representationalism

Recently, Chalmers has developed a view about the nature of mental qualities that he claims is influenced deeply by Shoemaker's accounts, but that departs in some crucial ways from them (Chalmers 2004, p. 367). Chalmers calls this view Fregean representationalism, a position which he maintains can accommodate both quality inversion and the central idea of representationalism that mental qualities are intimately related to representational properties.

§3.3.1 Fregean Representationalism Introduced

The heart of Chalmers's account of inversion is the notion of Fregean content, which Chalmers claims captures Gottlob Frege's (1892/1980) classic distinction between the sense and reference of an expression (for a similar account, see, e.g., Thompson 2009). Frege famously observed that one can believe that the sentence "Hesperus is bright" is true, and yet also simultaneously believe that the sentence "Phosphorus is bright" is false, even though both 'Hesperus' and 'Phosphorus' refer to the same object—namely, the planet Venus. Roughly, Frege claimed that though the expressions 'Hesperus' and 'Phosphorus' share the same referent they exhibit different senses, which he claimed contain or involve distinct modes of presentation of that referent.

There is much debate about how to best understand Frege's notion of modes of presentation, but Chalmers claim is that we can understand a mode of presentation to be a descriptive condition on an expression's reference. The mode of presentation of 'Hesperus', for

example, might be something like *the object usually visible at a certain point in the evening sky* (Chalmers 2004, p. 362). This description is a condition on the expression's reference insofar as any sentence involving 'Hesperus' is true only if the sentence concerns the object usually visible at a certain point in the evening sky. That is, a mode of presentation is a kind of truth condition; just as Frege held that sense determines reference, Chalmers holds that an expression's Fregean content in turn determines what it refers to.

Importantly, Chalmers claims that not only do linguistic expressions exhibit Fregean contents, but qualitative states such as perceptions do as well. Moreover, Chalmers argues that it is reasonable to think that each mental quality is equivalent to the property of having some Fregean content, where this Fregean content is something like *the property that normally causes experiences that exhibit this mental quality in me*. Chalmers writes,

> Let us say that Fregean representationalism is the thesis that phenomenal properties are equivalent to certain (pure or impure) Fregean representational properties…. For example, one might propose that phenomenal redness is equivalent to the property of having a certain Fregean content (in the appropriate phenomenal way), where this Fregean content involves a mode of presentation such as the property that normally causes experiences of phenomenal redness (2004, p. 364).

This characterization of the view is a bit concentrated and requires a measure of unpacking. In order to understand Chalmers's proposal, and thereby expose some difficulties for it, it will be helpful to explain how Chalmers's account is intended to account for cases of quality inversion. Unlike the versions of standard representationalism discussed in chapter 2 (versions which are not compatible with quality inversion), Fregean representationalism does not hold that the contents of perceptions are individuated solely in terms of the external properties that those states represent there to be. Rather, Fregean representationalism holds that inversion is possible because the contents of perceptions are individuated also in terms of the modes of presentation by which those external properties are represented.

To put Chalmers's view in context, let us recall that Shoemaker maintains that the possibility of quality inversion is explained by Jack and Jill representing the same perceptible object as exhibiting two distinct external appearance properties. Chalmers, by contrast, claims that the possibility of inversion is explained by Jack and Jill representing the same ordinary perceptible properties but via distinct mental properties. Chalmers's idea is roughly akin to Frege's conception of distinct expressions referring to the same object. Frege argues that 'Hesperus' and 'Phosphorus' include distinct modes of presentation of the same referent Venus. Chalmers similarly argues that two perceptions can represent the same perceptible property by via different perceptual modes of presentation.

But like Shoemaker's accounts, Chalmers's proposal holds that Fregean contents are modes of presentation that represent perceptible properties as the typical causes of mental qualities. On Chalmers's view, a perception that exhibits a given mental quality is veridical just in case one perceives there to be whatever property it is that normally causes perceptions with that mental quality. Since the properties responsible for causing mental qualities are the ordinary perceptible properties such as colors, perceptions that exhibit Fregean contents represent perceptible properties. And because Fregean contents are characterized in terms of typical causes, perceptions that exhibit distinct Fregean contents can represent the same perceptible property and perceptions that exhibit the same Fregean content can represent distinct perceptible properties. Chalmers writes,

> Of course, this mode of presentation may pick out different properties in different environments, or in creatures with different perceptual systems. In me, it might pick out a certain surface spectral reflectance; in a subject in a different environment or with a different perceptual system, it might pick out a different property entirely. But this is just what one would expect (2004, p. 365).

For example, if Jack's perception exhibits mental red, Chalmers claims that Jack's mental quality determines that his perception has the Fregean content, roughly, *that there is the property that normally causes perceptions with mental red in me*. And suppose that the property that normally causes perceptions with mental red in Jack is perceptible red. In that case, Jack's perception's Fregean content determines that it represents there to be perceptible red. So if Jack perceives something perceptible red, then Jack's perception is veridical; otherwise, Jack's perception is misrepresentational.

But suppose Jack's inverted twin Jill perceives something perceptible red, a perception which exhibits mental green. In that case, Jill's mental quality determines that her perception has the Fregean content *that there is the property that normally causes perceptions with mental green in me*. If the property that normally causes perceptions with mental green in Jill is perceptible red, then her perception is veridical. In short, Chalmers claims that both Jack and Jill can veridically represent the same perceptible property and yet exhibit distinct mental qualities, because those distinct qualities determine different Fregean contents.

§3.3.2 Problems with Fregean Representationalism

Despite its ingenuity, Chalmer's view faces many difficulties.[75] First, Chalmers claims not only that perception represents properties such as perceptible colors, but also that perception constitutively involves reference to itself. That is, Chalmers argues that a perception's Fregean contents represent there to be the causes of its mental qualities. But one might observe that this surely is not how things appear introspectively. Common sense suggests both that perception

[75] For example, as Boyd Millar (2012) recently argues, Fregean representationalism struggles to explain phenomena such as color constancy. But I will not pursue this kind of objection here.

represents colors as being mind-independent properties of objects and that perception does not comment on how those properties cause the perception's own properties.

In reply to this challenge, Chalmers claims that the fact that perception does not seem to refer to itself is explained by the fact that Fregean contents are nonconceptual. He writes,

> One might object that the sort of modes of presentation I have been discussing over-intellectualize the contents of experience. When one attends to a red ball, one does not usually conceive of it as the cause of one's experience, or as possessing properties that normally cause that sort of experience. But this point is compatible with what I have said. The Fregean contents I have appealed to may very often be nonconceptual contents: to have a state with these contents, a subject need not deploy a concept with those contents. So a subject's visual experience can have a mode of presentation along the lines of the object causing this experience without the subject deploying the concept of causation, or of the experience (2004, p. 368, emphasis his).

Chalmers does not offer much elaboration upon what it is for content to be nonconceptual, though I will defend a version of the view that mental qualities are nonconceptual contents in chapter 5. My only point here is that Chalmers's claim regarding Fregean content is puzzling, especially in light of what he has to say about Fregean modes of presentation in a related passage concerning the transparency intuition:

> One might wonder whether the Fregean view is compatible with the oft-noted transparency of experiences. As I construe it, the central datum of transparency is that when we attempt to introspect the qualities of our experiences (e.g. phenomenal redness), we do so by attending to the qualities of external objects (e.g. redness). In effect, we look "through" the phenomenal property…. Nevertheless, the mode of presentation exists, and one can become introspectively aware of it (2004, p. 368).

On the one hand, Chalmers holds that we can be introspectively aware of our Fregean modes of presentation. But if so, then Chalmers should agree that we can be introspectively aware that our perceptions represent their causes. Yet on the other hand, Chalmers maintains that this is not how introspection reveals perception. Ultimately, it is unclear how describing those contents' as nonconceptual would bear upon the issue of whether or not perception introspectively appears to represents its causes. A great deal rides on how Chalmers understands

these nonconceptual contents, but without such an account, Chalmers's stand on how introspection reveals perception seems untenable.

But even if Fregean content has the form that Chalmers hypothesizes that it does, I'll argue that his account nonetheless fails as an explanation of mental qualities. First, consider that Chalmers holds that a mental quality is not only equivalent to the property of having a Fregean content, but equivalent to having that content in the appropriate phenomenal way. The purpose of this qualification, as Chalmers acknowledges, is to account for the fact that one could have a belief that exhibits a Fregean content and yet not be a perceptual state. Chalmers writes,

> It may well be that a belief could attribute a property under this mode of presentation, however, so one needs to at least specify a perceptual manner of presentation. As a result, it seems plausible that phenomenal redness is equivalent to the representational property of perceptually phenomenally representing the relevant Fregean content (2004, p. 365).

But this qualification raises a potentially grave difficulty for Chalmers's view. What does representing something in an appropriate phenomenal way amount to? If Chalmers thinks that a state's being qualitative involves standing is some sort of perceptual attitude towards a Fregean content, then it is clear that his account suffers from the same kind of problem that I have raised for the versions of standard representationalism discussed in chapter 2. Chalmers requires an informative account of what it is to represent a content in a phenomenal way, which he arguably cannot give. If, for example, the perceptual attitudes somehow simply involve mental qualities, it would seem that this view smuggles in mental qualities without explaining them.

Second, Chalmers claims that a mental quality is equivalent to the property of having a certain Fregean content—but what does being equivalent amount to? There may be a variety of ways to elucidate this claim, but I'll argue that regardless of how it is understood, it is unclear how mental qualities determine or place constraints on Fregean contents. One might think that

Chalmers has in mind that each mental quality is identical to the property of exhibiting some Fregean content in an appropriate phenomenal way. But this interpretation seems to be ruled out by Chalmers' own remarks:

> It is highly plausible that two phenomenally identical experiences will have the same Fregean content. This will certainly be the case if the Fregean content of perceptual experiences works as above, by requiring a certain relation to the relevant phenomenal property [namely, by mentioning the quality in the content]…. This suggests that a phenomenal property entails the corresponding Fregean representational property…. As for the converse, it seems plausible that any mental state that visually phenomenally attributes a property under this mode of presentation will itself be phenomenally red (2004, p. 365).

If mental qualities were identical to the property of having a Fregean content, then there would be no question whether two qualitatively identical perceptions exhibit the same Fregean content. But Chalmers seems to think that there is such a question. Thus, instead, it seems that Chalmers holds that a mental quality's equivalence to the property of having a particular Fregean content amounts to the content's supervening on the quality. On this view, one's exhibiting mental red entails that one exhibits the Fregean content *that there is the property that normally causes mental red in me* just in case one exhibits mental red. But Chalmers offers no good reason to think that contents supervene in this way. Chalmers does assert that it is plausible that two states with the same mental quality will have the same Fregean content, but he provides no argument supporting this assertion. Moreover, even if a Fregean content mentions a particular quality, that fact does not entail that one must exhibit that quality. What we need is a good reason to think that mental green cannot determine the Fregean content that there is the property that normally causes perceptions with mental red. Or why one couldn't exhibit a mental quality but no Fregean content or exhibit a Fregean content but no mental quality? Chalmers offers no reason to think that one's mental qualities and Fregean contents could not come apart.

Without some reason to think that Fregean contents supervene on mental qualities, Chalmers's account at best entails that mental qualities accompany Fregean contents, but that mental qualities are not themselves representational. The representational work is done by the Fregean contents—and mental qualities' only function is to be mentioned by those contents. But that leaves open questions as to the nature of those qualities and their relationship to the contents. As Joseph Levine observes,

> By making phenomenal character a sense—a mode of presentation—and not merely a functional property, one legitimates the idea that something genuinely representational is going on here. However, how is the phenomenal character supposed to be a mode of presentation? How is it presenting the reference, the surface spectral reflectance? One idea would be to characterize the sense descriptively as "the distal property responsible for the reddish phenomenal quality of this experience"…. The problem here, though, is that what's doing the mode of presentation work is the description, not the phenomenal character itself. Notice that the phenomenal character is mentioned in the description, and its role there is purely as a property instantiated by the visual state (or the subject). The only representational role it plays is satisfying the effect role in the descriptive condition that picks out the distal cause (2010, p. 214).

Since the representational apparatus that is posited to explain inversion only makes mention of mental qualities, it is not clear why a state must exhibit any quality at all. Perhaps Chalmers holds that my having a sensation of red just is my representing there to be a thing that normally causes mental red in me, but I in fact exhibit no mental red. This would entail the perhaps undesirable consequence that we systematically misrepresent things, but Chalmers does not offer any explanation of the relationship between content and qualities that would rule this possibility out.

If mental qualities are separable from what they represent, then we arguably have no independent grasp on their identities. Since Chalmers's view collapses into the view that there is no substantive connection between a mental quality and what a state represents there to be, he thereby imports mental qualities without giving any informative account of them. In other

words, Fregean representationalism is at best a form of weak representationalism, according to which every qualitative state exhibits Fregean content, but also exhibits nonrepresentational mental qualities of the sort that traditional qualia theorist posit.[76] Speaking of a view like Chalmers's, Shoemaker observes that "there is of course a danger that by making the mode of presentation a feature of the perceptual experience this view sneaks in by the back door the qualia which standard representationalism is anxious to reject" (2003, p. 260). Indeed, Chalmers's theory does not seem to be an account of mental qualities at all.

§3.4 Block on Mental Paint

Block has long been a critic of representationalism and a staunch defender of the traditional qualia view (see, e.g., Block 1990, Block 2007a, Block 2010). But his theory of mental qualities merits consideration not only because it presents a major alternative to the version of strong representationalism that I will develop in chapter 5, but also because Block argues that some mental qualities do play a role in representing perceptible properties. Block claims we can distinguish the following sorts of mental properties of the experience of seeing a red tomato:

> 1. The representational content of an experience. I am currently looking at a tomato and my experience represents the tomato as red. 2. Mental properties of the experience that represent the redness of the tomato. This is mental paint. According to me, the phenomenal character of the experience is such a mental property (2003, pp. 541-542).

Block proposes that so-called mental paint is representational because he wants to explain the kind of inversion wherein Jack and Jill both veridically perceive a red tomato as being

[76] I take this to the spirit of Kriegel's remark: "Sometimes one finds in the relevant literature mention of a "Fregean representationalism" according to which phenomenal character is understood not in terms of what a conscious experience represents but in terms of how it represents it (i.e., in terms of the experience's Fregean mode of presentation)…. But, as I use the terms, Fregean representationalism is a rubber duck (it is not really a kind of representationalism)" (2009, pp. 72-73, fn. 17).

perceptible red, but the tomato looks red to Jack and looks green to Jill. As noted above, it is doubtless that the word 'looks', as well as related expressions such as 'appears' and 'seems', are ambiguous in several ways. And as also noted above, Shoemaker (1982, p. 365) claims that such expressions have an intentional reading, which is why he claims that Jack and Jill must be representing the tomato as exhibiting appearance colors distinct from its perceptible color. But Block claims that such a use of 'seems' instead tracks states' mental qualities.[77] He writes,

> In what Shoemaker (1981 [*sic*]) calls the intentional sense of 'looks the same', the chips look the same (in respect of color) to Jack and Jill just in case both of their perceptual experiences represent it as having the same color.... I recognize another sense of 'looks the same', (the qualitative or phenomenal sense), a sense geared to phenomenal character (2003, pp. 560-561).

According to this reading of 'seems' and related expressions, the tomato's seeming red to Jack amounts to Jack's having an experience that exhibits mental red when Jack sees the tomato. Thus Block's account of inversion is that the difference between Jack and Jill is not a difference in what is represented—that is, a difference at the level of representational content—because they both veridically perceive the tomato to be perceptible red. Rather, Block suggests that there is a difference in the qualities with which Jack and Jill represent the tomato's perceptible color. Block writes,

> If things we both call 'red' (and think of as red) look to you the way things we both call 'green' (and think of as green) look to me, then we have phenomenal qualities that represent ripe tomatoes as red, but do so very differently: your mental paint represents differently than mine. Further, we can have experiences that are the same in respect of mental paint, but represent differently, for example when I am looking at a red thing and you are looking at a green thing (2010, p. 55).

[77] Block's conception of 'seems' here echoes what Chisholm (1957, chapter 4) called the noncomparative use of such expressions. This is the use operative in sentences such as "The apple seems red to Jack." Following Frank Jackson (1977, p. 30), this has come to be known as the phenomenal use.

For example, a tomato will look red to Jack because his perception exhibits a certain kind of mental paint—mental red—in virtue of which it represents the tomato's red color. By contrast, that tomato will look green to Jill because her perception represents the tomato's red color by exhibiting a distinct kind of mental paint, mental green. Additionally, Block imagines that things that are green will look green to Jack, but look red to Jill. That is, the very same kind of mental paint might at one time or for one individual represent some perceptible property, but at another time or for another individual represent different perceptible properties. When a state exhibits mental paint, it thereby also exhibits some representational content. Since Jack's mental red functions to represent perceptible red and since Jill's mental green functions to represent perceptible red, Jack's states that exhibit mental red and Jill's states that exhibit mental green both also exhibit the content *that there is red*. In other words, as Block claims, "[t]he relation between mental paint and representational content is many-many" (2010, p. 55).

Block's use of the expression 'mental paint' is no accident. Block is conceiving of mental paint on analogy with the physical kind of paint one might find on a map. Consider, for example, a spot of blue paint on a map that represents Hawaii. The spot of blue paint on the map need not represent Hawaii because, naturally, a spot of red paint might just as easily have represented it. A spot of perceptible red and a spot of perceptible blue both represent the same location on a map, but they do so in different ways. Moreover, there is clearly a way in which neither spot of perceptible color is representational because the perceptible colors need not represent anything at all outside of the context of the map. As Loar puts it,

> Like paint on canvas, qualia are individuated, on this view, independently of their representational or referential properties, and—again like paint—individuated independently even of their purporting to represent, independently of their having representational properties even in Brentano's sense (2003, p. 78).

Mental paint is nonrepresentational insofar as it is only in virtue of mental colors playing a role in people's behavioral repertoires that qualitative states represent anything at all. Mental qualities enjoy, as it were, independent lives beyond their roles in representing perceptible properties. Block writes that "the phenomenal character of color experience, for example, could be said to be nonrepresentational in that the identity of that phenomenal character is not given by its representational content" (2003, p. 542). But the view counts as a form of weak representationalism insofar as qualitative states exhibit representational content determined by the role its mental paint plays.

Recently, Block has characterized his view in terms that echo Frege's distinction between the sense and reference of an expression.[78] As noted above, there is much debate about how to best understand Frege's notions. But Block suggests that we can think of mental qualities as analogous to the modes of presentation of expressions. On Block's view, just as one can refer to Venus using the word 'Hesperus' and yet another can refer to Venus via 'Phosphorus', so too can two people represent the same perceptible property via distinct mental qualities. Mental paint is in that way part of the content of perception. Block writes,

> [T]hink of the representational contents of perception as individuated both in terms of reference and mode of presentation, where qualia are taken as modes of presentation of those referential contents as suggested by Tyler Burge (2003). Nothing hangs on how we use the term 'representational content' (and I will continue to use the term to mean color-representing purely referential content). The important point is that if normal people can be inverted with respect to one another, there is an aspect of color experience that cannot be captured in ordinary language in terms of properties of objects (2007b, p. 85).[79]

[78] In this way, Block's view resembles Chalmers's account of Fregean representationalism. But Chalmers's characterization of modes of presentation is more specific than Block's and his view seems to depart substantially from Chalmers's account in ways that will become clear shortly. But to the extent that there are similarities between them, my critique of Chalmers's view also applies to Block's.

[79] There is, however, one obvious difference between Frege's and Block's conceptions of modes of presentation. On Frege's view, sense determines reference insofar as two expressions cannot have the same sense and yet exhibit distinct referents. Block, by contrast, urges that inverted

On Block's view, the fact that we use similar expressions to describe perceptible properties and mental qualities is accidental. There is nothing about mental red that connects it to perceptible red because for some individuals mental red represent perceptible red whereas for others it may represent perceptible green. Our ordinary ways of describing mental qualities invariably fall short of capturing the phenomena. But this is consistent with Block's cautionary point about mental qualities that I referenced in chapter 1: That if you have to ask about the nature of mental qualities, you simply will never understand what they are.

§3.4.1 Mental Qualities that Represent Per Se

But even if we accept these problematic details of Block's view for now, several questions about the view emerge. Block clearly holds that cases of intramodal inversion—that is, inversions within a given sensory modality—are possible. Mental red can represent perceptible red for me, though it might represent perceptible green for you. But it is not clear whether or not Block thinks that cases of intermodal inversion (wherein the mental qualities of one sensory modality represent the perceptible properties of a different modality) are possible. For example, does Block think that it is possible that the mental quality of the smell of roses might represent perceptible green for some creatures and the quality of the sensation of the sound of trumpets might represent perceptible triangles for others?

---

individuals' experiences can exhibit the same mode of presentation—the same mental quality—and yet represent different perceptible properties. So it is unclear how to unpack Block's proposal and, unfortunately, he does not say much more about this way of understanding modes of presentation. Indeed, even within the philosophy of language, the notion of a mode of presentation is somewhat obscure (see, e.g., Zalta 2001). But whatever the details of Block's claims about modes of presentation, the basic shape of Block's account is that different people can exhibit distinct mental qualities, yet those mental qualities can function to represent the same perceptible property, and the same mental quality can function to represent different mental qualities for different people.

Consideration of the analogy of the physical paint on a map may be helpful here. There is nothing about the spot of blue paint that makes it a candidate for representing Hawaii or that determines that it represent Hawaii rather than another geographic entity or something else altogether. The paint is nonrepresentational, but can be used in some context to represent anything that can be represented. We can hold up a sample of blue paint and say or think that it represents economic injustice, and it represents economic injustice. Within some context, anything can play the role of representing anything else.

One might think, then, that Block regards mental paint as simply any mental property that represents a perceptible property for some creature. So long as a property is a property of mental states, it is capable of functioning as mental paint by representing some perceptible property. On this view, it follows that both intramodal and intermodal inversions are possible. Though Block is not explicit regarding whether or not he regards intermodal inversion as possible, I'll shortly explore some grounds for denying this possibility.

Whatever his position on intermodal inversion may be, Block clearly denies that any mental quality can represent any perceptible property. Block further distinguishes the mental qualities that can represent many different kinds of perceptible properties from those mental qualities that represent, as he says, per se—that is, a mental quality that only represents a particular perceptible property. Using 'P-content' (short for 'phenomenal content') to refer to mental qualities, Block writes,

> Consider the P-content of seeing a square compared to the P-content of seeing a circle. These P-contents allow one to see that the squares are packable together without gaps, whereas the circles are not. Also, the squares have a small number of axes of symmetry, but the circles have a large number. These examples show that P-content is representational, but they also show something stronger and more interesting… what the examples show is that P-contents represent per se (1995, p. 227).

Block argues that perceptible squares exhibit certain geometrical features, such as what he calls packability, which perceptible circles do not. One can pack several perceptible squares together so that there is no space between them, but one cannot do so with perceptible circles. Moreover, Block claims that reflection alone on the mental qualities of sensations of those perceptible shapes reveals the fact that the perceptible shapes possess these geometrical properties. Block writes that "P-contents alone allow one to see such facts" (1995, p. 227).

As Block recognizes, these considerations entail that a given spatial mental quality cannot represent distinct perceptible properties. If one can tell introspectively that one's sensation is of a packable perceptible shape, then one's mental shape quality cannot represent a nonpackable perceptible shape because the discrepancy between the way the perceptible shape is and the way one represents that perceptible shape to be would be behaviorally detectable. Unlike mental color qualities, it would seem a particular mental shape quality can only represent a particular perceptible shape.

If mental qualities represent only particular perceptible properties, then those mental qualities cannot be undetectably inverted. Because mental shape qualities represent per se, a mental square quality cannot represent a perceptible circle and a mental circle quality cannot represent a perceptible square. Block claims that the distinction between mental qualities that represent particular perceptible properties and those that do not roughly tracks the so-called primary/secondary quality distinction that I discussed in chapter 1 (e.g., 1995, p. 227).

In what follows, I'll argue that Block does not have good reasons to distinguish these varieties of mental qualities. I'll argue that the kinds of considerations that Block adduces in favor of thinking that the qualities of sensations of shape represent only those shapes will apply equally to the qualities of sensations of color. So Block must conclude either that all qualities

represent particular qualities or that any mental quality can represent any perceptible property. If the latter is the case, however, Block does not offer any informative account of mental qualities.

§3.4.2 Problems with Block's Account

Let us consider again the question of intermodal inversion. Block might dismiss the possibility of intermodal inversions on the following grounds: Just as mere reflection on one's sensation of a shape reveals which perceptible shape the sensation represents, so too can mere reflection reveal that one is having a sensation of a perceptible color and not, for example, the sensation of a perceptible sound. But I'll argue that considerations of this sort reveal that intramodal inversions are impossible.

Block asserts that reflection on one's sensation of a shape reveals which perceptible shape the sensation represents because reflection on that mental shape discloses, for example, whether that perceptible shape is packable. But many distinct kinds of perceptible shapes are packable. For example, both certain perceptible rectangles and certain perceptible triangles can be arranged to fit together without gaps, but perceptible rectangles are packable with one another in a different way than perceptible triangles are packable with one another. Likewise, different perceptible shapes are not packable for different reasons. Different kinds of ovals fit together in ways that create different kinds of gaps. We might call a perceptible shape's kind of packability (or lack thereof) its packability profile. Each perceptible shape must have its own particular packability profile if, as Block maintains, reflection on one's mental shape reveals which particular perceptible shape one is representing.

But just as it is impossible to invert spatial mental qualities because they represent perceptible shapes as having uniquely identifying geometrical properties, so too, I argue, do

perceptible colors exhibit certain features that render it impossible to invert the mental colors that represent them. It is well-known that the perceptible colors vary at least along the dimensions of hue, saturation and brightness (see, e.g., Hardin 1993, p. 26). The very same hue—for example, perceptible red—can be saturated or unsaturated to varying degrees as well as bright or dull to varying degrees. The dimensions of hue, saturation, and brightness generate a color-ordering space, wherein each location in the system defines some ratio of the color dimensions (see, e.g., Hardin 1993, pp. 116-121). The fact that saturated perceptible red resembles saturated perceptible orange more than saturated perceptible blue is reflected in the fact that saturated perceptible red is closer in the space to saturated perceptible orange than it is to saturated perceptible blue. Call this space the quality space of the perceptible colors.[80]

It is also well-known that the quality space of the perceptible colors is asymmetrical. For example, after taking into account small differences between human discriminatory capacities, the most saturated perceptible yellow that a person can discriminate is brighter than the most saturated perceptible blue that a person can discriminate (for more examples of asymmetries in the quality space of the perceptible colors, see, e.g., Byrne 2010, §2.3.1). Each perceptible hue thus possesses what I'll call a saturation-brightness profile unique to that hue.

This aspect of the quality space of perceptible colors opens Block's account to a variant of a much-discussed objection to the conceivability of quality inversion. Because the quality space of the perceptible colors is asymmetrical, many have argued that it is not possible for mental colors to be inverted in a behaviorally indistinguishable way (e.g., Clark 1985, p. 438-439; Tye 1995, pp. 203-204). If, for example, Jack were to represent perceptible blue via mental

[80] I will discuss these and other features of quality spaces in greater detail in chapter 5.

blue, but Jill were to represent perceptible blue via mental yellow, then Jill would judge certain blue objects as brighter than Jack would judge them to be.

There is a related problem for Block's account of mental paint. Block argues that mental square represents a perceptible shape with a particular packability profile and so cannot represent shapes that are not packable, such as perceptible circles. If this is the case, however, I argue that the dimensions of saturation and brightness are to perceptible hue as packability is to perceptible shape. In other words, it is likewise reasonable to think that mental blue represents a perceptible color with a certain saturation-brightness profile and so cannot represent a perceptible color with a different profile. Moreover, just as reflection on one's mental shape alone allows one to see that a certain perceptible shape exhibits certain features, reflection on one's mental color alone allows one to see that a certain perceptible color has certain features. Thus the same considerations that Block adduces to rule out the inversion of the purportedly primary qualities equally undermine the inversion of the purportedly secondary qualities.

Block might reply that it is easy to see on reflection on one's mental shape whether the perceptible shape it represents can fit without gaps with other perceptible shapes. By contrast, Block might continue, it is difficult or impossible to see whether a particular hue can exhibit a certain level of brightness simply by reflecting on one's mental color. But it is only because of our familiarity with perceptible shapes that we are able on reflection to appreciate that a particular perceptible shape can fit with other shapes without gaps. If one were not familiar with a particular perceptible shape, then one would arguably not know on reflection whether or not the shape is packable.[81] Likewise, if one is very familiar with perceptible colors and their

[81] Cf. Pautz (ms.), p. 24, fn. 33. Pautz notes that his three-year old daughter does not know that a perceptible shape is packable, even though she presumably can introspect the qualities of her spatial perceptions.

relations to one another, one might know, simply on the basis of reflection on one's mental color, that a perceptible color has a particular relative capacity for saturation or brightness. For example, it does seem fairly obvious from introspection alone that mental yellow corresponds to a perceptible hue that can be brighter than any perceptible blue. So there is no good reason to think that mental shapes reveal the features of perceptible shapes whereas mental colors reveal nothing about the features of perceptible colors.

Instead, Block might reply that while the perceptible colors do in fact exhibit unique combinations of hue, saturation, and brightness, this state of affairs need not have been the case. Speaking of the evidence that the space of perceptible colors is asymmetrical, Block writes that "even if such experiments do show some asymmetries, there are possible creatures—maybe genetically engineered versions of humans—in whom the asymmetries are ironed out" (2003, p. 536). Block's reply to the evidence that the actual quality space of perceptible colors is asymmetrical is that it is at least possible that there be creatures that have symmetrical spaces (see also, e.g., Shoemaker 1982, p. 367).

But if Block is willing to "iron out" the asymmetries in the space of color qualities, it is not clear why the same sorts of considerations could not apply to other sorts of mental qualities such as the spatial qualities. Perceptible shapes resemble and differ from one another in such a way that generates a quality space of perceptible shapes (see, e.g., Meehan 2001, Rosenthal 2005). So, for instance, perceptible circles are more similar to perceptible ellipses than they are to perceptible squares. A quality space of the perceptible shapes would reflect these similarities and differences among the perceptible shapes, including in respect of their packability. The packable perceptible shapes would be closer to one another within the space of perceptible shapes than they would be to the nonpackable perceptible shapes. We can then ask: If the quality

space of perceptible shapes is, or at least conceivably could be, symmetrical, why can't mental square represent a perceptible circle?

Since Block cannot explain why some qualities represent particular perceptible properties whereas others do not, he has three options: He can hold that any mental quality can represent any perceptible property, that mental qualities are not representational at all, or that each mental quality represents only a particular perceptible property.

Suppose that Block maintains that any mental quality can represent any perceptible property. On this view, mental paint is just whatever mental quality happens to be playing a representational role for a creature. One problem with this view is that there is as yet no explanation of why a given mental quality happens to be playing its representational role. Recall, Block is thinking of mental paint on analogy with physical paint on a map. In the case of a map, however, we have an independent explanation of why a given spot of perceptible color represents a particular location. A spot of blue paint does not represent Hawaii simply because it is blue, but because it is being used by creatures with intentional states to represent Hawaii.[82] It is because of the blue paint's role in the behavioral economy of a creature that it plays the role of representing Hawaii, though almost anything could equally play that representational role.

Unfortunately, this sort of explanation does not clearly apply to the case of mental paint. It does not seem as though a creature uses its mental qualities to represent perceptible properties in the same way that it uses a spot of physical paint to represent geographic entities on a map. Nor does it seem as though creatures can opt to use a mental quality to represent properties distinct from the property it does represent. Block might reply that any instance of mental paint

[82] To adapt Searle's distinction discussed in chapter 1, the spot of blue paint is only derivatively representational; the spot's representational character depends on and is explained by its relations to genuinely representational states.

just happens to play the functional role of representing a given perceptible property. But without an explanation of how mental paint plays these functional roles, it is not clear that this reply amounts to any informative explanation of mental qualities.

Similarly, if mental paint is simply whatever mental quality happens to be playing a representational role for a creature, then it is unclear why mental paint can only represent perceptible properties, as opposed to any other kind of property. After all, there are no restrictions on what a spot of blue paint can represent. Likewise, what could restrict what a mental quality represents? If mental red can represent perceptible green or a perceptible square, what prevents it from representing nonperceptible things such the property of being a dog or economic injustice? Without an explanation of the mechanism that explains why mental qualities represent at all, there is arguably no limit to what mental qualities can represent.

The biggest problem with Block's view is that, lacking an account of why certain mental qualities represent what they do, Block's story about mental paint adds nothing to our understanding of mental qualities. In the case of the spot of perceptible paint on the map, not only do we have an independent explanation of why a given spot of perceptible color represents a particular location, but we also have an independent grasp of the identity of the perceptible color doing the representing. We have ways of individuating perceptible colors independent of their occasional roles representing locations on maps.

But in the case of mental qualities, if they are capable of being undetectably inverted, it is not clear what independent understanding we have of them. Indeed, this is the central problem that afflicts the traditional qualia view. On that view, it is not that Jack and Jill represent the same perceptible color via distinct mental colors; rather, it is that Jack and Jill represent the same perceptible color and simply happen to exhibit distinct mental colors. If mental qualities are not

representational at all or can represent anything, then we require another explanation of them. But it is not clear that Block is offering any explanation of them. Thus, either Block's complex analysis of mental paint collapses into the traditional qualia view, or each mental quality represents only a particular perceptible property. I will defend a version of the latter view in chapter 5, though this view has the consequence that quality inversion is impossible.

§3.5 Conclusions

I have argued that none of these complex versions of strong and weak representationalism that seek to accommodate the possibility of quality inversion provide informative accounts of mental qualities. In the end, however, this sort of consequence may not trouble these theorists. After all, Chalmers and Block hold that there is a hard problem about mental qualities and that consciousness is the only route to understanding them.[83] Chalmers himself explicitly acknowledges that his version of representationalism is "doubly nonreductive, with phenomenal elements involved in both the manner of representation and the representational content" (2004, p. 370).

But if Chalmers's view is a nonreductive account of mental qualities, then we require some way to account for our knowledge of them, even from the first-person perspective. The most promising explanation of our knowledge of mental qualities is the so-called phenomenal-concepts approach, and both Block and Chalmers have endorsed versions of this view. In the next chapter, I'll argue that the major accounts of phenomenal concepts too are unworkable. The reasonable conclusion, then, is that no informative theory of mental qualities can accommodate

[83] Shoemaker, by contrast, denies there is a hard problem about mental qualities. As noted above, Shoemaker provides a theory of mental qualities, which I will explore in detail in chapter 5. I'll conclude, however, that Shoemaker's theory fails to account for mental qualities.

quality inversion. In chapter 5, I'll develop a theory of mental qualities according to which quality inversion is impossible, but I will also argue that this consequence does not present an obstacle to a theory of mental qualities.

CHAPTER 4: PHENOMENAL CONCEPTS

§4.1 Introduction

I argued in chapter 3 that versions of representationalism that attempt to accommodate the possibility of quality inversion are unworkable and ultimately collapse into the traditional qualia view. One might conclude, then, that some version of the traditional qualia view is true. But, as noted in chapter 1, if quality inversion is possible, then we cannot have knowledge of others' qualia from the third-person perspective. If it were possible that someone else's perception of a blueberry might exhibit mental yellow instead of mental blue—and if there is no way to determine which quality it does exhibit—then the only genuine access we have to mental qualities is via our own first-person awareness of them. Thus the traditional qualia theorist owes an account of how we know about mental qualities—even from the first-person perspective.

If one endorses the traditional qualia view, perhaps the most promising account of our knowledge of such nonrepresentational qualia from the first-person perspective is what is known as the phenomenal-concepts approach. According to this approach, we can explain the role that mental qualities play in our psychological functioning through positing a distinctive class of concepts known as phenomenal concepts ("p-concepts") (e.g., Loar 1990/1997, Hill and McLaughlin 1999, Perry 2001, Papineau 2002, Chalmers 2003, Tye 2003b, Carruthers 2004, Aydede and Güzeldere 2005, Block 2007c, Levin 2007, and Balog 2012).[84]

P-concepts are unusual insofar as they purportedly refer to one's own current mental qualities in a direct and nondescriptive way. Moreover, p-concepts putatively refer to one's own

[84] Not all proponents of the p-concepts approach are traditional qualia theorists. For example, Tye (2003b) endorses a version of the view, though he is a representationalist (it is worth noting, however, that Tye (2011) recently abandons the approach). Nevertheless, my goal in this chapter is to argue that this p-concept approach is the most promising way to save the traditional qualia view—and yet is nonetheless fails to do so.

qualities in a way that does not merely seem direct, but is direct because the concepts somehow depend on or are partially constituted by the qualities they refer to. According to the p-concepts approach, one can acquire a particularly secure sort of knowledge of our own mental qualities—what is often referred to as phenomenal knowledge—by forming thoughts about them that literally involve the qualities themselves.

These theories of p-concepts present an alternative understanding of the role mental qualities play in our psychological lives to the role that I will develop in chapter 5, according to which mental qualities are themselves a class of representations. According to the p-concepts approach, mental qualities function in our mental economies in a way that does not require that they are representational. One might think that since p-concepts conceptualize mental qualities and since concepts are representations, the p-concepts view holds that mental qualities are representational. But just as deploying the concept DOG to think about a dog does not require that the dog is itself a representation, it may be that we conceptualize mental qualities with p-concepts, but that those qualities are not themselves representational.[85]

As I'll discuss in more detail shortly, the p-concepts approach was engineered primarily to explain how physicalism is compatible with various intuitions about mental qualities, such as the intuition that quality inversion seems possible. If such scenarios are in fact not possible, then we require some mechanism to explain how our access to mental qualities would give rise to the impression that physicalism is false. In chapter 5, I'll argue that scenarios such as quality inversion are plainly not possible. Thus insofar as the p-concepts approach has been invoked to

[85] Indeed, as I'll discuss shortly, most posit p-concepts in an attempt to show how physicalism is compatible with various antimaterialist intuitions. Since these intuitions are typically taken to show that mental qualities are, by themselves, nonrepresentational (e.g., the conceivability of undetectable inversion), most proponents of p-concepts hold that mental qualities are by themselves nonrepresentational (e.g., Block 2007c).

explain a questionable set of givens at the outset, the complex machinery that they employ is unmotivated. I'll argue, moreover, that there are independent reasons to think that none of the well-known accounts of p-concepts work. In particular, I'll argue that none successfully explain how p-concepts refer to mental qualities. Since I suspect that no viable theory of p-concepts can be developed, the traditional qualia views that depend on such accounts do not constitute credible alternatives to the theory of mental qualities that I'll defend in chapter 5.

§4.1.1 The Basics of the P-Concepts Approach

We can think about mental qualities in a variety of ways. For example, I can think about the quality of a pain that I am currently having. But I may also think about the quality of a pain that someone else has or about the quality of a pain that I experienced two days ago. As Chalmers (2003, pp. 223-226) argues, the p-concepts I use to refer to others' or my own future or past mental qualities refer in relational ways—that is, via descriptions such as "the quality that Jack is having now" or "the quality that is typically caused by seeing red things." The p-concepts at issue, by contrast, putatively refer to mental qualities in a direct and nondescriptive way. Chalmers calls these p-concepts direct p-concepts (2003, p. 235). From now on, my use of 'p-concept' will refer to these sorts of directly referential and nondescriptive p-concepts, unless otherwise specified.

The idea that some concepts can refer directly to their referents finds its roots in prominent work in the philosophy of language dating from the second half of the 20th century. One of the leading theories of how expressions such as names, natural-kind terms, and demonstratives refer is the so-called direct reference theory, pioneered by Kripke (1980) and others (see also, e.g., Kripke 1971, Putnam 1975, Devitt 1981, Kaplan 1989). Kripke famously

argues that names such as ‘Aristotle’ and natural-kind terms such as ‘water’ do not refer via associated descriptions. Descriptions associated with a name only serve to “fix the reference” of such expressions (e.g., 1980, p. 55). Instead, Kripke suggests what he calls a “better picture” (e.g., 1980, p. 97), according to which such expressions refer directly by standing in appropriate causal relations to their referents.

Proponents of the p-concepts approach similarly hold that p-concepts share an intimate causal connection to their referents because they somehow constitutively depend upon the qualities to which they refer. Several theorists have offered different ways to understand how p-concepts are constituted by their referents and I will discuss a few of these proposals shortly. Expressing this roughly, one can token a particular p-concept only if one tokens as well the quality to which it refers. A thought involving a p-concept would look something like the following: THE QUALITY OF MY CURRENT EXPERIENCE IS THIS (cf. Chalmers 2003, p. 236), where the concept THIS is a p-concept that constitutively involves the quality to which it refers.

Most hold that p-concepts’ being so constituted by the mental qualities to which they refer explains why p-concepts refer in a direct and nondescriptive way. Discussing these sorts of accounts of p-concepts generally, Katalin Balog writes,

> More precisely, on this view, every concept token applied to current experience is constituted by a current token phenomenal experience, and—on most versions of the constitutional account—this fact is crucial in determining the reference of the concept (2009, p. 306).

Importantly, because p-concepts refer in this fashion, p-concepts putatively do not conceptualize their referents in terms, for example, of their neural bases (if they have such a basis), of how they relate to one’s behavior, or of their psychological connections to other states. Rather, as

Chalmers puts it, this sort of concept simply "characterizes the phenomenal quality as the phenomenal quality that it is" (2003, p. 226). Likewise, David Papineau writes,

> When we use phenomenal concepts, we think of mental properties, not as items in the material world, but in terms of what they are like. Consider what happens when the dentist's drill slips and hits the nerve in your tooth. You can think of this materially, in terms of nerve messages, brain activity, bodily flinching, facial grimaces, and so on. Or you can think of it in terms of what it would be like, of how it would feel if it happened to you (2002, p. 48).

At this point, one might observe parenthetically that it is quite unclear exactly what it is for a p-concept to pick out a quality as the particular quality that it is. I'll argue shortly that there is in fact no way to explain how it is that a p-concept could refer in this fashion. But for now I will assume that we have a sufficient grasp of what is indicated by the notion of p-concept to proceed with the basics of that approach.

That p-concepts refer in a direct and nondescriptive way purports to explain a variety of features of our awareness of mental qualities. For example, most people would agree that it often seems quite difficult to describe our mental qualities from the first-person perspective. If I smell a flower, it may seem to me that I am only able to describe the qualities of my experience in meager ways. I might say that the quality is sweet or rosy or floral, but it is not clear that I can say much more than that. If it is true that p-concepts refer in a direct and nondescriptive way, then it is unsurprising that mental qualities are relatively indescribable in this fashion—we are not aware of them in a manner that would enable rich descriptions of them. Likewise, it is because of this intimate relationship of the quality we have thoughts about and the way we have thoughts about it that our access to mental qualities appears to be so direct and so special. Our awareness of our own mental qualities seems direct because it is direct.

Indeed, it is often assumed, especially by those working within a Cartesian framework, that we have <u>incorrigible access</u> to our own mental qualities. And according to many proponents

of the p-concepts approach, p-concepts can explain why it is that we have this kind of incorrigible access, at least they explain it to some extent (see, e.g., Balog 2009, p. 293). Since a thought that one's current experience exhibits some mental quality putatively involves a p-concept constituted by that very quality of one's current experience, it would seem that such thoughts cannot be false. Calling beliefs about the qualities of one's current experiences direct phenomenal beliefs, Chalmers writes,

> A direct phenomenal concept by its nature picks out instances of an underlying demonstrated phenomenal quality, and a direct phenomenal belief identifies the referent of that concept with the very demonstrated quality (or predicates the concept of the very experience that instantiated the quality), so its truth is guaranteed (2003, p. 242).[86]

As noted before, p-concepts are often posited to explain how physicalism is compatible with various antiphysicalist thought experiments, such as Frank Jackson's (1982) famous scenario involving Mary the super-scientist. Jackson's argument centers on Mary, who is imagined to have been brought up in a room wherein the color of all of the objects, surfaces, and so forth are some achromatic shade (that is, shades on the black-gray-white spectrum). But Mary is also imagined to be a super-scientist, who knows all of the facts about anything physical, including facts about how perceptible colors, vision, the brain, and so on work. Upon her release from her achromatic room, Mary comes across a red tomato. According to Jackson, Mary learns about something new—namely, she learns about mental red, which to date she had never exhibited. Thus, Jackson concludes, facts about mental qualities are not facts about anything physical. In other words, physicalism is false.

According to what Daniel Stoljar (2005) has called the p-concept strategy, physicalism is safe from these sorts of thought experiments. According to the strategy, when Mary sees the red

[86] As Chalmers recognizes, this at best establishes a very limited incorrigibility claim, according to which only one's beliefs about one's occurrent conscious qualities are incorrigible. Beliefs about one's future or even recently past qualities are corrigible (2003, p. 242).

tomato for the first time, she does not learn any new facts. Rather, proponents of the p-concepts strategy argue that after her release Mary is able to think about her mental quality, which she could previously think about, in a new way. It is not that Mary could not think about mental qualities while in the achromatic room. She could think, for example, that she would exhibit mental red if she were to be released and see a red tomato. But prior to release, Mary only possessed concepts of her mental qualities that characterized those qualities descriptively. Having never seen anything red, Mary could not yet exhibit mental red and so could not deploy a nondescriptive p-concept of her quality.

Despite appearances to the contrary, the strategy maintains that physicalism is true because p-concepts such as MENTAL RED and concepts of physical things such as SUCH AND SUCH PATTERN OF NEURAL FIRINGS are coreferential, though the latter (but not the former) can be tokened in the absence of a mental quality. In other words, there are two kinds of concepts, not two kinds of entities. P-concepts corefer to concepts of physical or functional things (presumably neural states), but since they are nondescriptive, this identity is neither immediately apparent nor easily explained. As Loar claims, "Phenomenal concepts are conceptually irreducible in this sense: they neither… imply, nor are implied by, physical-functional concepts" (1997, p. 597). Though the scenario involving Mary does not threaten physicalism, the p-concept strategy seems to offer an explanation of why it seems that it does.[87]

[87] I note that Chalmers (2003, 2007) himself also posits p-concepts, but nonetheless argues that physicalism is false. Chalmers puts forward p-concepts to explain how it is that we can know, for example, that we are not philosophical zombies, even though he claims that it is conceivable that there are such beings in other possible worlds. But Chalmers also argues that p-concepts and physical concepts cannot corefer and that dualism is nonetheless true. For a reply to Chalmers on behalf of the phenomenal concept strategy, see, e.g., Carruthers and Veillet 2007. My interest here, however, is not in whether positing p-concepts can save physicalism, but whether there is a viable account of p-concepts. For reasons that I discuss in chapter 5, I do not think that we need to posit p-concepts, nor does physicalism need to be rescued from this sort of

§4.2 Loar's Recognitional P-Concept Account

There are various ways to examine the proposal that we conceptualize our mental qualities via special p-concepts. Loar (1990/1997) provides what is perhaps the canonical presentation of the p-concepts approach. Loar holds that p-concepts are a subclass of what he calls recognitional concepts, which are "type demonstratives… grounded in dispositions to classify, by way of perceptual discriminations, certain objects, events, situations" (1997, p. 600).[88] A type demonstrative is a demonstrative concept of the form THAT F. When I see a dog, I might think to myself "That dog is brown." In that case, I deploy a recognitional concept of it, the concept THAT DOG, which I possess if I am disposed to recognize it as being of that type (i.e., dog) as opposed to some other type.

Loar's recognitional p-concepts similarly demonstratively pick out a particular occurrent mental quality as being of the type that it is. A p-concept has the form THAT QUALITY and a phenomenal thought might have the form MY CURRENT EXPERIENCE HAS THAT QUALITY. The concept is thus something of a mental analogue to the linguistic expression 'that F', which can, on distinct occasions, be used to refer to any token of the F type.

Loar's recognitional p-concepts are, however, very different from other recognitional concepts and are not reducible to any other kind of physical, functional, or demonstrative concept. P-concepts, unlike these other sorts of concepts, involve the qualities that they refer to as their modes of presentation. Loar writes,

> We might say that a phenomenal concept has as its mode of presentation the very phenomenal quality that it picks out. We might also say that phenomenal concepts have "token modes of presentation" that are noncontingently tied to the phenomenal qualities

---

antiphysicalist argument. So I will not evaluate Chalmers's arguments here; I will argue that the p-concepts approach fails for other reasons.

[88] For similar recognitional theories of p-concepts, see, e.g., Tye 2003b, Carruthers 2004, Levin 2007.

> to which those concepts point: particular cramp feelings and images can focus one's conception of the phenomenal quality of cramp feeling (1997, p. 604).

As noted in chapter 3, the notion of a mode of presentation is a bit obscure, but Papineau clarifies Loar's proposal thus:

> [T]he idea that phenomenal properties can provide their 'own modes of presentation' may simply mean that they do not have to be picked out via some other contingently connected property they possess. There is only one property in play where a phenomenal concept refers to a phenomenal property: namely, the phenomenal property itself. No further property mediates between referring concept and referent (2002, p. 104).

In other words, we can roughly understand Loar's proposal in the general way that I described above: A particular p-concept token refers to a mental quality token as being of some type of mental quality by constitutively involving the token to which it refers. The tokening of a quality in the right manner somehow enables the quality to refer to its type in a direct fashion. Importantly, because p-concepts are demonstrative, they do not in any way describe the qualities to which they refer. This latter aspect is what gives the impression that p-concepts do not corefer to physical or functional concepts. But this, Loar maintains, is a mistake; p-concepts refer to physical properties.

Despite its ingenuity, several problems beset Loar's account. First, there is considerable debate regarding whether or not recognitional concepts even exist. Jerry Fodor (1998b) argues, for example, that there are no recognitional concepts because concepts are compositional, but recognitional capacities are not. As Fodor has long argued (e.g., Fodor 1975), concepts are compositional in the following way: If one possesses two concepts and basic conceptual capacities, one is able to combine those two concepts into a complex concept. And, on Fodor's view, the semantic properties of the complex concept are fully determined by the semantic properties of the component concepts. So, to borrow Fodor's example, if one possesses the concepts PET and FISH as well as basic conceptual capacities, then one thereby possesses the

concept PET FISH. The complex concept PET FISH is constructed from the separate concepts PET and FISH. And, Fodor claims, the meaning of PET FISH is determined solely by the meaning of its constituent concepts.

If any concepts merit the label 'recognitional concept', then the concept PET and the concept FISH surely present as excellent candidates. But, as Fodor argues, one may possess the concept PET and the concept FISH and be able to recognize pets and to recognize fish as such, but yet not be able to recognize pet fish as such. Fodor concludes that one's capacity to recognize that to which a concept refers plays no role in that concept's semantics. The recognitional capacity is not preserved when one combines component concepts to form a complex whole. Fodor concludes therefore that there is no such thing as a recognitional concept—and thus no p-concepts in Loar's sense.

Some have replied to Fodor in defense of recognitional concepts. Terry Horgan (1998) and François Recanati (2002) both argue that, while the semantic properties of a complex concept are inherited from the semantic properties of its constituent concepts, recognitional properties are not heritable in this fashion. It is worth noting that neither Horgan nor Recanati offer any reasons to think that recognitional properties would not be heritable and it is not obvious that they would not be. Without resolving the issue altogether, the point can be made that it is not at all obvious that there are such entities as recognitional concepts.

But a very damaging difficulty for Loar's account is that there is no way to specify the mechanism by which a p-concept refers to the relevant mental quality. On Loar's view, p-concepts are demonstrative concepts that refer to distinct token qualities on different occasions as being of the same type. Consider by analogy the linguistic demonstrative 'that'. Linguistic demonstratives do refer to different things in different contexts of use. On one occasion, I may

use ‘that’ to refer to a cat, but on another occasion use ‘that’ to refer to a dog. But in order to understand any given use of a linguistic demonstrative, there must be some context of use that both determines what it is that it refers to, and enables others to understand that reference. If there were no such context of use, then it would not be clear that the demonstrative would refer to anything at all.

One might think that p-concepts do not require further explanation of why they refer to the relevant qualities. After all, Loar claims that the relevant qualities function as the modes of presentation of the p-concepts and in that way constitute the concepts. But even if a quality constitutes a concept, it does not entail that the concept refers to the quality. For one thing, it could be that other aspects of the concept are responsible for the concept’s referring. But even if the quality’s constituting the concept somehow explains why the concept refers, the concept is constituted by the quality however that quality is described. And since Loar is a physicalist, each quality is identical to some physical property. Loar is clear, however, that though p-concepts corefer to physical properties, they do not refer to them as such. P-concepts pick out qualities as qualities, not as physical properties. Thus some further story must be told illuminating why p-concepts refer to qualities as qualities, even if those qualities constitute them. Discussing Loar’s account, Papineau similarly argues that it “betrays loose thinking about reference to suppose that concepts automatically refer to any properties that are involved in their deployment” (2002, p. 106).

Loar says alarmingly little about what accounts for p-concepts’ referring. But there are many approaches to explaining how linguistic demonstratives refer, which may provide avenues

to explain how these demonstrative p-concepts refer. I'll consider several of these proposals in turn.[89]

If one rejects, as does Loar, that the reference of an expression or concept is determined descriptively, then the most obvious way to explain reference proceeds in terms of a Kripkean account, according to which reference is secured by expressions standing in appropriate causal connections to their referents. And some have argued that the appropriate causal relation that determines the reference of linguistic demonstratives is secured by perception. This view is propounded by, among others, Gareth Evans, who writes,

> We are now in a position to answer the question what makes demonstrative identification of spatially located material objects possible. In the ordinary perceptual situation, not only will there be an information-link between subject and object, but also the subject will know… upon the basis of that link, where the object is (1982, p. 170).[90]

On this view, if I say, for example, "That is red," the referent of 'that' is determined by the fact that I am currently perceiving whatever is red—I, as it were, perceptually point to the referent of 'that' and thereby stand in a reference-determining causal connection to the referent. But whether or not this perceptual model works in the case of linguistic demonstratives, it is not clear that it can explain how p-concepts demonstratively refer because we arguably do not perceive our mental qualities. There are many ways in which our awareness of our mental properties is quite distinct from our awareness of external mind-independent perceptible properties. For example, there is no dedicated sensory organ involved in our awareness of mental qualities,

[89] The following discussion is much influenced by Josh Weisberg's (ms, chapter 3), wherein he examines several possible ways to explain the reference of p-concepts as it relates to Loar (1990/1997) and Papineau (2002).

[90] And this view has been more recently championed by Devitt, who asks: "In virtue of what is a thought a singular thought about a certain object in mind…. I argued that the thought is simply one that is grounded in the object by a certain perceptual causal process (Devitt 1974: 191-2)" (Devitt 2004, p. 290).

though perception typically involves such an organ (see, e.g., Shoemaker 1994a or Rosenthal 2004 for many arguments against perceptual models of our awareness of mental qualities).

One might reply that, whatever sort of process introspection may be, the appropriate reference-determining causal link is established by introspection. Tye, who also endorses a similar recognitional account of p-concepts, proposes this. Tye writes,

> [P-concept C] refers to a phenomenal quality Q via C's being the concept that is exercised in an introspective act of awareness by person P if, and only if, under normal conditions of introspection, Q is tokened in P's current experience and because Q is tokened (2003b, p. 7).

Tye's idea is that a p-concept refers to the quality it does because one forms that p-concept if and only if one introspects that quality. But without a clear account of introspection, it is not obvious how this proposal works. There are, of course, several theories of introspection currently available. There are higher-order thought accounts (e.g., Rosenthal 2005), higher-order perception accounts (e.g., Armstrong 1968, Lycan 1996), mindreading-module accounts (e.g., Carruthers 2011), inferentialist accounts (e.g., Dretske 1995), and others (for a survey, see, e.g., Schwitzgebel 2010 or Smithies and Stoljar 2012). I cannot settle here what type of process introspection might be. But whatever kind of process introspection is, if it is nonperceptual, then it is arguable that introspection involves descriptions of mental qualities. If introspection is a thought-like process or proceeds simply by forming thoughts about our mental qualities, then introspection almost certainly involves descriptive content. Moreover, since introspection classifies mental states as being certain ways, it is reasonable to think it is a descriptive process. But if the p-concept involved in introspection refers via some descriptive mechanism, then a major motivation for the p-concept approach is undercut.

One might reply, however, that introspection does not work descriptively, but merely by establishing a relevant causal link to a quality.[91] But even if the causal connection is established by introspection, a mere causal connection to a mental quality is ill-suited to explain how p-concepts refer. If all there is to explain why a demonstrative refers to an object is a causal connection between them, it is unclear how finely grained the reference of demonstratives can be. After all, causal connections are cheap insofar as there are countless causal connections between any two particular things. For example, there is a causal connection between a demonstrative and an object just in case there is also a causal connection between the demonstrative and the object's surface. The problem for the p-concept theorist is that p-concepts do not refer to mental qualities in any way; rather, p-concepts specifically refer to qualities as the qualities that they are. So what the p-concept theorist needs is an account of the causal connection that explains why the concept refers to the quality as the quality it is, and not as something else such as its neural basis.

One might hope that the resources to provide such an account might be found in Kripke's own discussion of the reference of names and natural-kind terms. Recall that Kripke argues that the reference of names and natural-kind terms is not determined by descriptions.[92] Roughly, Kripke instead proposes that a linguistic item such as 'Aristotle' refers to a certain person because a user of that item designates that person (here, Aristotle) as the referent of the name by a process that Kripke calls an "initial baptism" and there is then a successive causal process of passing that name from one user of the name to another (e.g., 1980, p. 96). Each successive user

[91] I thank Devitt (personal communication) for proposing this possibility.
[92] I will not canvass any of Kripke's well-known arguments against the so-called descriptive theory of reference here. Whether or not Kripke's proposal works in the case of certain linguistic items, I'll argue that it is unclear how this view of reference can work in the case of p-concepts.

of the name, by intending to use the name to refer to the same referent as her predecessor, succeeds in referring to that referent because the name is causally linked to its referent via the process of passing along the name from one user to another.

As Kripke himself notes, something must explain why the initial baptism succeeds in linking the expression to the relevant thing. Though he does not provide a complete account of this process, Kripke suggests that the initial baptism is "explained in terms of either fixing of a reference by a description, or ostension (if ostension is not to be subsumed under the other category)" (1980, p. 97). So, for example, the name 'Aristotle' might have come to refer to Aristotle because Aristotle's parents pointed at him and named him Aristotle. Or, for example, the name 'Aristotle' might have come to refer to him because his parents had the descriptive thought that the baby in the cradle will bear the name 'Aristotle'. On Kripke's picture, after this kind of initial descriptive or ostensive encounter with a referent fixes it as the referent of an expression, successive users of an expression can refer directly to a thing without deploying descriptive content to refer to it.

But whatever the credibility of this account with regard to names, natural-kind terms, and linguistic demonstratives, it is unclear how this account could succeed in the case of p-concepts. The proponent of p-concepts must, of course, explain how it is that mental qualities are initially baptized in a way that enables successive usages of the p-concepts to refer directly. If the initial baptism is some sort of ostensive process, since we cannot literally point to our mental qualities, the only reasonable candidates are perceptual or introspective acts. But, as I've argued, we arguably do not perceive our mental qualities. And, since what we are currently seeking is an explanation of what makes it the case that introspection refers to mental qualities in the first place, we cannot appeal to introspection to explain the initial baptism. The other option is that

the initial baptism is some sort of descriptive act, but this option is also not open to the proponent of p-concepts. Proponents of p-concepts maintain that no descriptions can capture the nature of mental qualities. Indeed, this idea is the main reason for thinking that p-concepts refer directly. If we can refer to mental qualities in a descriptive way, however, then the central motivation for the p-concept approach is undermined.[93]

Another suggestion is that p-concepts refer to mental qualities in a way that mirrors a conversational process, which, following Michael Devitt and Kim Sterelny (1999, p. 58), I'll call reference-borrowing. If my conversational partner says, "That is red," using 'that' to refer to a tomato, I might in turn reply "Yes, and that is round." In this case, what might explain the fact that my use of 'that' succeeds in referring to the tomato is that I am intending to use the same referent as my interlocutor. I therefore borrow the reference from my partner, so that we both succeed in referring to the same thing. One might think that I likewise refer to my mental quality by somehow borrowing the reference of my p-concept from former usages of the concept.

But reference-borrowing is a public, conversational process that involves multiple users of the same linguistic item. The use of p-concepts, by contrast, is by hypothesis a private affair insofar as it involves an individual's own mental qualities. In the case of a p-concept, the concept has not been used previously by others to denote the given quality, nor has the concept been passed along in such a manner that could secure its reference. So it is uncertain at best how to apply the reference-borrowing notion to p-concepts.

[93] Kripke does acknowledge that there may be other ways that initial baptisms could take place. After offering the characterization above, Kripke writes, "Perhaps there are other possibilities for initial baptisms" (1980, p. 97). But without some other account of how the reference of p-concepts is determined, any account at best remains incomplete and at worst stipulates that p-concepts refer directly.

Instead, one might argue that the reference of p-concepts is determined by relevant background beliefs. If we are discussing a red tomato sitting on a desk and I say, "That is red," the referent of 'that' might be determined by what is salient to the conversation. What we are talking about is the tomato. We both have background beliefs about what we are discussing, which determine the reference of the demonstrative. Likewise, one might think that a p-concept refers to a particular mental quality because one holds a relevant background belief about it. I might believe, for example, that I currently have been seeing a red tomato. Thus my thought about the quality of my perception refers to the quality it does because of that belief.

But this proposal will not do the job for the p-concept theorist because p-concepts by hypothesis refer directly, whereas background beliefs would again determine the referent of a p-concept in a descriptive way. The content of my p-concept would thus be something like the descriptive concept THE QUALITY OF MY PERCEPTION OF THE RED TOMATO. While it is reasonable to think that we refer to our mental qualities by forming descriptions of them in this manner, this proposal would be tantamount to denying that p-concepts refer in any direct way. Thus it is unclear how p-concepts directly refer on Loar's account.

§4.3 The Quotational Model

Instead of explaining p-concepts as a subclass of so-called recognitional concepts as Loar does, many theorists have attempted to explain p-concepts on the model of linguistic quotation (e.g., Papineau 2002, Block 2007c, Balog 2012). Once more, proponents of the p-concepts approach maintain that p-concepts refer directly and in a nondescriptive way to the qualities to which they refer because they somehow depend on those concepts. And many have claimed that linguistic

quotations work in something like this way too—and so we can understand how p-concepts operate by understanding how linguistic quotation works.

I note at the onset that there is a bit of difficulty in characterizing linguistic quotation. Linguistic quotation is more or less a syntactic operation wherein an expression is designated in such a way that the quotation mentions but does not use the quoted expression (see, e.g., Cappelen and Lepore 2007, chapter 1; for more on the use/mention distinction, see, e.g., Quine 1940, p. 24). Typically quotation involves placing a quoted expression within quotation marks.[94] So, for example, the expression 'cat' refers not to cats, but refers to the word that designates cats. 'Cat' thus mentions, but does not use, the word for cats. This explains why the sentence "'Cat' has three letters" is an intelligible sentence in English in addition to being true, whereas the sentence "Cat has three letters" is unintelligible.

It would seem thus that when one quotes a word, the quoted word literally constitutes the quotation. In the case of linguistic quotation marks, Papineau claims that it would seem that:

> The referring term incorporates the things referred to, and thereby forms a compound which refers to that thing. Thus, ordinary quotation marks can be viewed as forming a frame, which, when filled by a word, yields a term for that word (2002, p. 117).

Some theorists have argued that a p-concept likewise refers to a quality by involving that quality in a way somehow like quotation. On this view, p-concepts are a sort of mental analogue to linguistic quotation, which use mental qualities to mention themselves. Papineau pioneered this mental quotation model of p-concepts. He writes,

> [P]henomenal concepts are compound terms, formed by entering some state of perceptual classification or re-creation into the frame provided by a general experience operator 'the experience: - - -'…. [M]y phenomenal concepts involve a frame… and, when this frame is filled by an experience, the whole then refers to that experience (2002, pp. 116-117).

[94] I say 'typically' because there are ways to quote an expression that do not involve quotation marks. As, for example, Marga Reimer (1996, p. 135) observes, we can quote an expression in various ways, such as by putting it in italics (cf. Cappelen and Lepore 2007, p. 36).

The basic idea is that a p-concept of a quality of my current experience, say mental red, consists in a frame, what Papineau calls the experience operator 'the experience: - - -', which is filled in by my mental red. And Papineau emphasizes that not just any experience will suffice to fill in the experience operator. Papineau is firm that one cannot, of course, use another person's experience to fill the frame—it must be one's own experience (2002, p. 123). When the experience operator is filled in through this process, it refers to the quality of one's own experience, mental red. As Papineau puts it, "roughly speaking, we refer to a certain experience by producing an example of it" (2002, p. 116).

But a major question that faces any account of quotation, mental or otherwise, is this: What explains why a quotation refers to the particular expression that it does? Much like Loar, Papineau holds that mental qualities, by participating in mental quotations, serve as the modes of presentation of themselves (2002, pp. 103-104). But as noted before, Papineau acknowledges that this alone does not explain why such modes of presentation refer to what they do. However, like Loar and other proponents of p-concepts, Papineau cannot claim that p-concepts refer to the qualities they quote via descriptive means. Papineau insists that "phenomenal concepts refer directly, and not by description" (2002, p. 87).

Thus Papineau seeks to offer a different account of how mental quotations refer in a nondescriptive way to their quoted mental qualities. There are several accounts of linguistic quotation on offer, and Papineau adopts something akin to Davidson's demonstrative theory of quotation (see also, e.g., Cappelen and Lepore 2007). Davidson writes that "[linguistic] quotation marks…help refer to a shape by pointing out something that has it….The singular term is the quotation marks, which may be read 'the expression a token of which is here'" (1979, p. 90). Papineau likewise writes,

> Linguistic quotation marks, after all, are a species of demonstrative construction: a use of quotation marks will refer to that word, whatever it is, that happens to be made salient by being placed within the quotation marks (2007, p. 121, emphasis his).

Before turning to some specific problems with Papineau's account, it is worth noting that there is a serious problem with any account of p-concepts that putatively explains them on the model of linguistic quotation. As noted previously, there is considerable debate concerning how linguistic quotation operates (cf. Cappelen and Lepore 2007). As a result, it is at best unclear whether or not expressions literally constitute quotations, and therefore it is likewise unclear how mental qualities literally constitute p-concept mental quotations in the way that the above accounts would require.

W.V. Quine, for instance, argues that quotations are syntactically atomistic (see also, e.g., Tarski 1933, p. 159). On Quine's view, a quotation does not literally involve its quoted expression, but is itself a novel expression that has nothing to do with the quoted expression except insofar as it refers to it. Quine writes, "From the standpoint of logical analysis each whole quotation must be regarded as a single word or sign, whose parts count for no more than serifs or syllables" (1940, p. 26).

This view of quotation has many virtues. For example, it is plain that one cannot substitute coreferential terms within a quotation and preserve truth values. For example, it does not follow from the fact that Twain is Clemens that one can substitute 'Twain' in a sentence involving 'Clemens'. The sentence "'Twain' has five letters" is true, whereas the sentence "'Clemens' has five letters" is false (cf. Cappelen and Lepore 2007, p. 100). Thus Quine colorfully writes of quotation, "To make substitution upon a personal name, within such a [quoted] context, would be no more justifiable than to make a substitution upon the term 'cat' within the context 'cattle'" (1961, p. 141). That is to say, a quoted expression is not a semantic

part of a quotation any more than ‘cat’ is a semantic part of ‘cattle’. This account is, of course, controversial and I cannot settle the issue here (see, e.g., Cappelen and Lepore 2007, chapter 9). My point is only that it is not at all obvious that a quoted mental quality is constituted by a mental quality.

The proponent of the quotational model of p-concepts might reply that the fact that linguistic quotations are not literally constituted by their quoted expressions does not show that mental quotations cannot be so constituted by mental qualities. That is, they might reply that the analogy with linguistic quotation is bound to go only so far; it might be that p-concepts nonetheless work the way, for example, Papineau hypothesizes.[95]

But as several others have pointed out, there are many problems with a demonstrative account of linguistic quotation (see, e.g., Cappelen and Lepore 2007, chapter 10). As noted above in connection with Loar’s recognitional account of p-concepts, a major problem for any account that relies on demonstratives is that it must explain why a demonstrative refers to what it does, as opposed to something else. If linguistic quotations are a kind of demonstrative, then we require some explanation of why a linguistic quotation demonstratively refers to the expression that is quoted, as opposed to some portion of the expression or to something else altogether. For example, why does ‘cat’ refer to the word for cat as opposed to ‘ca’ or to a red tomato (cf. Cappelen and Lepore 2012, §3.3.2, objection 1)?

Perhaps, as Herman Cappelen and Ernie Lepore (2007) argue, these problems for the demonstrative account of linguistic quotation can be overcome. But I do not see how this account can be successfully applied to the case of p-concepts, if they are a special kind of mental demonstratives. The fact that Papineau’s account differs from Loar’s insofar as p-concepts are a

[95] I thank Katalin Balog (personal communication) for noting this reply.

species of mental quotation seems to be of no help in answering this question. Indeed, since Papineau's account stands or falls with an explanation of mental demonstration as much as Loar's, it is unclear how these two accounts substantively differ.

Papineau does clarify his view by endorsing a teleosemantic account of reference in general. He writes that "[t]he referential value of [a] concept can then be equated with those items which it is the biological function of the concept to track" (2002, p. 113). On Papineau's view, a p-concept allegedly refers to the mental quality that it quotes because the concept has evolved the biological function of referring to that quality's type. But this emendation to the view does not help, because a parallel difficulty arises for it. What Papineau requires is an explanation of how p-concepts evolved these biological functions, but it is unclear how he can provide that story without bringing in descriptive content.

As we saw in connection with Kripke's account of initial baptisms, there are typically two ways that one can initially interact with something in order to designate it as the referent of an expression: Ostension or description. And one might likewise think that one of these mechanisms explains how it is that p-concepts acquire the biological function of referring to mental qualities in a nondescriptive way. But neither of these mechanisms is well-suited to explain this. First, it is unclear what kind of ostensive process could initially refer to mental qualities so that they could acquire such a referential function. Again, we arguably do not perceive our mental qualities, so perception cannot explain this process. One might think we initially refer to mental qualities via introspection, but that just pushes the question one step back. What explains how introspection refers to mental qualities? If, by contrast, we initially refer to our mental qualities by having descriptive thoughts about them, then not only is the central motivation for the p-concept approach undercut, but also it is unclear how p-concepts

could then acquire their further function of referring to themselves in a nondescriptive way. Without such an explanation, Papineau's account remains incomplete.

But one might pursue a different sort of account of the reference of mental quotations. Like Papineau, Block (2007c) does endorse a version of the quotational account of p-concepts. Because Block's account is quite similar to Papineau's, I will not explore its details. But I invoke it because Block offers an alternative account of the reference of p-concepts, which he calls a dispositionalist account. Block writes,

> There is an outstanding issue involving phenomenal concepts that I will raise briefly without attempting to resolve. What makes it the case that a token phenomenal property in a phenomenal concept serves as a token of one phenomenal type or property rather than another…. The answer on the dispositionalist view is that it depends on the subject's disposition to, for example, treat another experience as falling under the same concept. You are thinking that the experience is nice—but what will you count as another one of those? If only another experience as of Green126,731 will count as an experience of the same type, the phenomenal concept is maximally specific; if any bluish green experience will count as an experience of the same type, the concept is more general. If any experience of green will count as an experience of the same type, the concept is still more general (2007c, p. 283, n. 32).

Block's idea here seems to be that what makes a p-concept refer to the specific quality it does (opposed to a more or less determinate quality or a different quality altogether) is that when one deploys that p-concept one also exhibits a disposition to conceptualize other token qualities as tokens of that mental quality's type. But it is hard to see how Block's dispositionalist account constitutes a complete account of the reference of p-concepts, because it does not explain why one has the relevant reference-determining dispositions. In the attempt to explain this, it would seem that Block's account will run into an analogue of the problem raised for Papineau's teleofunctional account.

Again, the most reasonable ways to explain how we could form dispositions to recognize things in general proceed either in terms of our perceiving those things or by our having

descriptive thoughts about those things as being present. But if these are the mechanisms by which we come to be aware of things, it is not clear how to explain how we could acquire the kinds of dispositions to conceptualize mental qualities in a way that does not involve deploying descriptive content to refer to those qualities. But if the dispositions that purportedly determine the reference of p-concepts are dispositions to conceptualize mental qualities in a descriptive way, then it is unclear how those p-concepts could refer directly.

Balog, another proponent of the quotational model of p-concepts, offers yet another account of how their reference is secured. On Balog's view, a quotational p-concept refers to the quality it quotes because of the conceptual role of that concept. Balog claims that the reason why linguistic quotations refer to their referents is that users of quotations are disposed to accept instances of the following disquotational schemas (2012, p. 34):

L1: 'X' refers to X.

L2: ''X'' refers to 'X'.

These disquotational sentences characterize the conceptual role of linguistic quotation. But the sentences L1 and L2 are schematic insofar as they involve the variable 'X', which, when replaced with some token expression in English, form sentences regarding that to which a given quotation refers. And, Balog argues, it is because users of quotations are disposed to accept instances of these schemas that quotations such as 'cat' refer.

On Balog's view, p-concepts likewise refer to what they do because users of them are disposed to accept instances of similar disquotational schemas, which are the mental analogues of L1 and L2. Using Balog's terminology, suppose that 'experience x' refers to an experience x and 'REFERS-TO' refers to the concept of the reference relation. And Balog uses '*' as the mental analogues of quotation marks, which, instead of enclosing a linguistic expression, enclose

a mental item. Thus, on Balog's usage, '*experience x*' refers to a p-concept quoting the experience x. With this notation in hand, Balog argues that mental quotations refer in virtue of one's disposition to accept instances of the following thought schemas (2012, p. 34):

M1: *experience x* REFERS-TO experience x.

M2: **experience x** REFERS-TO *experience x*.

In this schema, 'M1' refers to a thought schema and when 'experience x' is replaced with a particular token experience, the thought is the thought that the p-concept that quotes that particular experience refers to that experience. Balog's account has some appeal because it is reasonable to think that the contents of many or even all concepts are determined by the conceptual roles they play (e.g., Harman 1987), although this is certainly controversial (e.g., Fodor and Lepore 1992). I cannot settle here the general issue of whether all concepts' contents are determined in a conceptual-role manner. But even if some concepts' contents are determined in a conceptual-role fashion, I will argue that Balog's account nonetheless fails for other reasons.

Balog, like most proponents of p-concepts, is working within the framework of the so-called language of thought hypothesis (e.g., Fodor 1975), wherein thoughts are composed of concepts that combine compositionally according to syntactic rules to form thoughts in a mental language often called Mentalese. On this view, concepts are to thoughts as words are to sentences. I will not dispute the language of thought model here. Instead, I will argue that, if one accepts the basic tenets of this model, then one cannot endorse Balog's account.

Importantly, as Balog herself recognizes, the right-hand expression of the thought schema M1 does not involve a concept of anything, but rather an experience. The token of 'experience x' in M1 on the left-hand side refers to a p-concept, but the token of it on the right-hand side that is not surrounded by the '*'s refers to the experience itself, not a concept of that experience.

And even if one held that representationalism were true, that would not entail that an experience is itself a concept. Indeed, Balog is clear that she does not regard experiences themselves as concepts. Thus Balog claims that people disposed to accept thoughts that are partially composed of concepts, but also composed of nonconceptual elements—namely, experiences. But Balog does not see this feature of her view as a problem. She writes,

> I do think experiences represent—but not conceptually so they are not in themselves concepts. This means that the expression on the right hand side of M1 has simply experiences, i.e., non-conceptual representations as instances. Some might object that the resulting thoughts will not be well-formed. However, I think there is a case to be made that plain—unquoted—experiences can be parts of thought under special circumstances (2012, p. 35).

It is unclear, however, how a nonconceptual element could fit with other concepts to produce well-formed or intelligible thoughts. Consider the case of linguistic communication. If I were to say "'Red' refers to," I would not have uttered an intelligible sentence, even if I were to then show you a patch of red. The most likely intelligible thing I could have been saying was "'Red' refers to red," where the second half of the sentence is filled not with a nonlinguistic item, but with the word 'red'. This is why, on a standard version of the language of thought framework, thoughts are composed of concepts and concepts alone. But even if Balog were to grant that thoughts involving p-concepts are sui generis, Balog's view faces another serious difficulty, one which is related to my objections to Loar's account. If the experience on the right-hand side of the M1 schema is unconceptualized, what explains why the p-concept on the left-hand side refers to it as the quality it is, as opposed to something else such as its neural basis or something else altogether?

Consider again the example of the patch of color and the sentence "'Red' refers to." Without any further information, it is not only unclear why 'Red' refers to the patch (instead of anything else), but also why it refers to it as a patch of a red, as opposed to as a picture, some

splotches of paint, and so on. We might explain that 'Red' refers to the patch of red as a patch of red because of one of the mechanisms that I proposed to explain how linguistic demonstratives work. When one reads the sentence, 'Red' might refer because one perceives the picture as a patch of red or because of certain background beliefs about its referent. But, again, these mechanisms cannot explain why p-concepts refer to qualities of experiences. Thus is it likewise unexplained how it is that users of p-concepts could be disposed to accept instances of schemas such as M1, let alone have thoughts that are instances of M1.

None of the major quotational accounts of p-concepts are thus capable of explaining how it is that p-concepts refer to their quoted mental qualities. Moreover, as argued above, it is unclear whether quotation works in this manner. In the absence of an instructive analogy for how to understand mental quotation, and in light of the failures of these accounts, it is reasonable to conclude that no quotational model of p-concepts can succeed.

§4.4 Papineau's Sensory-Template Account

In a recent paper, Papineau (2007) rejects his (2002) quotational model of p-concepts and instead offers an account of p-concepts which differs in some substantial ways from his initial proposal. Papineau drops the idea that p-concepts involve an experience operator and instead argues that p-concepts are a subclass of what he calls perceptual concepts. On Papineau's later view, perceptual concepts are "stored sensory templates" that refer to whatever they are templates of because they have acquired the functional role of doing so (2007, p. 114). For example, if I see a particular bird, I form a sensory template of that bird. My sensory template is a mental item that stores all of the perceptual information that I have acquired from that bird. And, Papineau argues, our sensory templates can be reactivated on other occasions. If I see the bird at other

times, then my sensory template of it is tokened again. Furthermore, my sensory template can accumulate additional sensory information about that bird on other occasions because it has the functional role of referring to that bird.

On this revised view, a p-concept is a sensory template of itself—a mental quality that has acquired the role of referring to qualities of its type (2007, p. 122). This new view shares many features with his old view: P-concepts are still constituted by their referents, refer to them in a direct and nondescriptive way, and so refer because of their teleofunctional role. And according to Papineau, the main similarity between his accounts is this:

> This model of phenomenal concepts as a species of perceptual concept retains one crucial feature from my earlier quotational-indexical model, namely, that phenomenal references to an experience will involve an instance of that experience, and in this sense will use that experience in order to mention it (2007, p. 123).

But Papineau asserts that his new position departs from the old one in critical ways. For example, Papineau is clear that, unlike his former account, his new account does not hold that these concepts are demonstrative. Papineau writes,

> I think it is quite wrong to classify perceptual concepts as demonstratives. If anything is definitive of demonstrative terms, it is surely that they display some species of characterlikeness. By this I mean that the referential value of the term depends on the context…. However, there is nothing characterlike about the kind of perceptual concept illustrated in the above examples (2007, p. 113).

As noted before, linguistic demonstratives refer to different things in different contexts of use. On one occasion, I may use ‘that’ to refer to a fish, but on another occasion use ‘that’ to refer to a dog. This is what it is for them to exhibit <u>characterlikeness</u>, to use Papineau’s expression (Papineau adapts the expression from Kaplan 1989, p. 505). But perceptual concepts do not, on Papineau’s new view, exhibit this context sensitivity. Whenever one deploys a perceptual concept, such as a sensory template of a bird, it always refers to its given referent—in this case,

the bird. On Papineau's revised view, a p-concept is a nondemonstrative concept that always refers to a particular type of mental quality (2007, p. 123).

Since Papineau's new account does not rely on demonstrative concepts, this view may appear to evade the objections that I have raised against his earlier account. But this is not the case. Like his old account, Papineau's new view requires an explanation of why mental qualities have developed the functional role of referring to themselves, and once again it is unclear how he can provide that explanation without bringing in descriptive content. Once more, perception and descriptive thought remain our two fundamental ways of understanding how we are aware of things, but neither of these processes is well-suited for explaining how mental qualities could acquire this kind of nondescriptive referential function.

Papineau's new view also faces other difficulties. It follows from this account, for example, that the fact that a p-concept is constituted by a mental quality is not what accounts for the concept's referring to the quality. Rather, it is the functional role of the concept that accounts for its referential properties. It is thus consistent with Papineau's new view that, as he himself acknowledges, "other states with a different or no phenomenology, but with the same cognitive function, would refer to the same experiences for the same reasons" (2007, p. 125).

Papineau, however, argues that there is good reason to think that mental qualities nonetheless do constitute p-concepts. Papineau claims that it makes sense that mental qualities would acquire the functional role of referring to themselves because they are arguably always apparent when such reference is taking place. Papineau writes, "Conscious experiences are excellent vehicles for thinking about those selfsame experiences simply because they are automatically present whenever their referents are" (2007, p. 126). But this hardly follows. No one thinks that the concept DOG has acquired the functional role to refer to itself, even if it were

true that the concept DOG is tokened every time the concept of that concept is tokened. So Papineau needs an explanation of why mental qualities would be special in this regard, and it does not seem he can provide such an account.

But even if Papineau were able to provide such an explanation, it is not obvious that it would be of any help. Since it is possible on Papineau's view for either mental red or a nonqualitative mental item to function to represent mental green, it seems that the original motivation for thinking that p-concepts are constituted by mental qualities in the first place has been lost. The fact that a concept happens to be constituted by the quality to which it refers appears to be irrelevant to whether or not the concept refers to a given quality. Moreover, if the quality need not be present for there to be a concept of the quality, it is not clear how the concepts could exhibit the various special referential properties that p-concepts are putatively thought to exhibit. How can mental red refer in a direct and nondescriptive way to mental green, if mental green is not tokened?

§4.5 Prinz's Mental Pointing

Though Prinz has been a vigorous critic of the p-concepts approach, in his (2007) he proposes an account of phenomenal knowledge that does have much in common with this approach. Prinz proposes that we refer to mental qualities not via p-concepts but by "mentally pointing" to them via nondescriptive demonstrative acts of top-down attention. Thus, like p-concept theorists, Prinz maintains that we refer to mental qualities in a direct way that requires the presence of the mental quality to which we refer. The view presents a new element, however, insofar as this mental pointing operation is not conceptual.

Prinz distinguishes mental pointing from a variety of mental operations. On one hand, Prinz claims that when one categorizes something, one brings concepts to bear on it that conceptualize the thing as being of some kind or another. By contrast, when one merely notices something, Prinz claims that one encodes that thing in working memory, whether or not it is conceptualized. And Prinz further claims that we can mentally point to things by attending to them in a top-down way, which need not involve concepts. Speaking of this mental pointing operation, Prinz writes,

> Let's consider the case of looking at a complex scene. It seems to be that there are two ways in which we can mentally point to an item in a scene. First, we can identify a particular region of space in which the item is located. Second, we can apply a perceptual concept…. If you are looking at a table setting, you might focus on the left and see what's there or you might look for the salad fork and focus on it, by using a high-level salad fork template. Notice that I used the word 'focus' in both cases. The phenomena I have just described is known in psychology as top-down selective attention…. I think mental pointing is achieved by top-down attention (2007, p. 194).

On Prinz's view, we achieve phenomenal knowledge of our mental qualities not by conceptualizing those qualities, but by attending to them in this pointing way. Of course, one might have thought that all knowledge is a conceptual affair. That is, to be in a state of knowing some content to be the case is to be in a state that exhibits the concepts that characterize that content. But Prinz argues that concepts of our phenomenal states are neither necessary nor sufficient for possessing phenomenal knowledge. They are not necessary, Prinz argues, because one can know what an experience is like, even if one does not possess the conceptual resources with which to describe that experience. He writes,

> Suppose you have never eaten guava before, but now, at this moment, have a first taste of guava in the form of a sauce on the fish you ordered. You can't identify the flavour you are tasting, and you might not recognize it if you tasted it again, but it might be true of you that you know what guava is like, nevertheless…. This shows that concepts used to categorize phenomenal states are not necessary for phenomenal knowledge (2007, p. 198).

But concepts are not sufficient for possessing phenomenal knowledge either, Prinz argues, because one can token concepts about one's experience even in the absence of consciousness. Thus those concepts need not involve the qualities to which they refer (2007, p. 199). Though categorization can take place outside of consciousness, Prinz maintains that mental pointing cannot do so because top-down attention is always conscious (2007, p. 200). Thus Prinz proposes that the best explanation of phenomenal knowledge is that it is nonconceptual and secured by mental pointing.

But there are several difficulties with Prinz's proposal. First, it is uncertain whether or not Prinz would subscribe to the p-concept theorists' claim that our awareness of mental qualities is an awareness of them as the qualities that they are. But Prinz does claim that an awareness of mental qualities is an awareness of "what it's like" (e.g., 2007, p. 196)—a statement which suggests that we are aware of them as qualities and not as, for example, their neural bases. If so, then Prinz's account faces a difficulty similar to the one raised for accounts which purported to guide the reference of p-concepts in terms of a mere causal relation. Again, what the p-concept theorist must offer is an explanation of the causal connection that explains why the concept refers to the mental quality as the quality it is, and not as anything else. But it is not obvious whether there could be such an account that does not bring in descriptive content. Likewise, since attention, on Prinz's view, is not conceptual, it is not clear that attention is the sort of process that can make us aware of mental qualities as anything, let alone as mental qualities and not under other descriptions.

Second, Prinz asserts that we are capable of deploying concepts of our experiences nonconsciously, so it cannot be the case that we use those concepts to acquire phenomenal knowledge. But as noted in previous chapters, though most assume that if mental qualities exist

they are always conscious, this assumption is questionable—and I'll argue that it is false in chapter 5. For example, there is substantial experimental evidence that one can perceptually discriminate stimuli without being aware that one perceives those stimuli (e.g., Breitmeyer et al 2004, Ro et al 2009). Since these so-called subliminal perceptions play the same kinds of perceptual roles as their conscious counterparts, it is reasonable to think that there are mental qualities that are not conscious. If mental qualities can occur in the absence of consciousness, then it is possible that we can use our nonconscious concepts to conceptualize them and form knowledge of them that is nonconscious.

Third, though Prinz assumes that top-down attention cannot occur outside of consciousness, this assumption is also questionable. There is growing evidence for top-down attention in the absence of consciousness (e.g., Koch and Tsuchiya 2007, van Boxtel et al 2010). Thus, even if all mental qualities did occur consciously, it is arguable that top-down attention is not sufficient for our being aware of them. But top-down attention is also arguably not necessary for awareness of conscious qualities either. We often think about or are aware of mental qualities without attending to them in a top-down way, such as in the case of the qualities on the peripheries of our conscious experience.[96]

Prinz (personal communication) argues that there is as of yet no evidence that the processes that occur nonconsciously or in the periphery encode information in working memory. But, on his view, it is crucial that mental pointing results in information being so encoded. Mental pointing just is the deployment of top-down attention and, on Prinz's view, attention is

[96] For brevity's sake, I will not discuss this evidence and I acknowledge that Prinz has disputed it on a variety of grounds (e.g., de Brigard & Prinz 2010). But a full discussion of whether attention is necessary and sufficient for consciousness (as he claims) is beyond the purview of this dissertation. I only note that there is mounting evidence that attention is neither necessary nor sufficient for consciousness and that the jury is still out on this deliberation.

"just a process by which perceptual information becomes available to working memory" (2007, p. 195). However, there is for example substantial evidence that stimuli that have been masked—and therefore rendered invisible to consciousness—can result in semantic priming (for review, see, e.g., Kouider and Dehaene 2007). And it is murky how nonconscious perceptions could yield such priming effects if the information from those states was not encoded in working memory.

Thus, whether or not we mentally point to our mental qualities via acts of top-down attention, these operations are arguably neither necessary nor sufficient to explain our awareness of mental qualities.

§4.6 Conclusions

I have argued that none of the major theories of p-concepts (or Prinz's related proposal) can explain how it is that they refer in a direct and nondescriptive way to mental qualities. And I suspect that no theory could explain this. But as noted in the Introduction, much of the motivation for thinking that our concepts of mental qualities must refer in this way depends on the traditional qualia theorists' assumption that mental qualities do not admit of third-person access or of informative description. If we assume this, then we need a mechanism to explain how it is that we do know about them, even from the first-person. But this assumption is unfounded and in the next chapter I'll argue that it is false. I'll then provide an informative account of mental qualities that eschews this questionable assumption. With this account in hand, it will be clear why we have no need for a special account of the concepts with which we conceptualize those qualities.

CHAPTER 5: HOLISTIC REPRESENTATIONALISM

§5.1 Introduction

In chapter 2, I argued that standard versions of representationalism that hold that qualitative content is conceptual or individuated in an atomistic way are problematic. Conceptualist versions of representationalism are unworkable primarily because they cannot offer an informative account of the difference between thought and perception and atomistic versions of representationalism face various other difficulties. But standard versions of representationalism are often rejected simply because they cannot accommodate phenomena such as the possibility of undetectable quality inversion. Instead of rejecting representationalism, some theorists have developed complex versions of the view that purport to accommodate this kind of possibility. But in chapter 3, I argued that these complex accounts fail to provide informative accounts of mental qualities and ultimately collapse into the traditional qualia view that there are nonrepresentational qualia that resist informative explanation. If, however, one accepts that traditional qualia view, then one must explain how we know about such nonrepresentational qualities, even from the first-person perspective. The most promising route to explaining this kind of phenomenal knowledge is the so-called p-concepts approach. But in chapter 4, I argued that none of the major versions of the p-concepts approach succeed in explaining how it is that p-concepts refer to mental qualities. It is thus reasonable to conclude that no informative account of mental qualities can accommodate scenarios such as quality inversion.

It might seem therefore that the traditional qualia view is true and we cannot put forward an informative account of mental qualities. But that conclusion would be premature. In this chapter, I'll develop and defend a better form of representationalism—a form of representationalism wherein perceptual content is individuated holistically. I'll begin by arguing

that the prevailing presuppositions that lead many to assume that mental qualities are nonrepresentational are unfounded.

§5.2 Two Approaches to Mental Qualities Revisited

As I argued in chapter 1, the traditional qualia view is largely driven by several related intuitions about mental qualities—such as the intuition that quality inversion is possible. These intuitions imply that mental qualities are intrinsic properties, at best contingently involved in the representation of perceptible properties if they play a representational role at all.

But as I also argued in chapter 1, these intuitions embody a certain assumption about mental qualities—namely, the assumption that mental qualities are only genuinely known from the first-person perspective in consciousness. Though this claim is rarely argued for and often simply assumed, many people working in the contemporary philosophy of mind do not seem to view this claim as an assumption but as a given. Most simply do not see any other way to understand what mental qualities might be, except as known from the first-person perspective. If, however, there were another way to understand mental qualities, a way that did not assume that they are only known from the first-person perspective, then the idea that they are only known in that way would be a theoretical claim in need of argument.

And there is just such an alternative approach—representationalism. According to representationalism, mental qualities are characterized in terms of their roles in the perceptual representation of perceptible properties. The representationalist approach does not explain mental qualities solely in terms of how they appear in consciousness. Rather, representationalism explains mental qualities in a way that permits third-person access to mental qualities as well. Representationalism does not, of course, deny that we often do have first-

person access to our own mental qualities. What the representationalist denies is that consciousness is the only genuine route to mental qualities. Thus a representationalist account of mental qualities must integrate with some of account of our first-person awareness of them. Because there is a view that holds that mental qualities can be known in a third-person way, the proponent of traditional qualia theory must provide good reason to think that consciousness is the only route to understanding mental qualities.

§5.2.1 Reasons to Reject the Traditional Qualia Theory[97]

There are, however, no good reasons to assume that mental qualities are only genuinely known from the first-person perspective in consciousness. One might think that the intuitions—such as the intuition that quality inversion is possible—support the claim that mental qualities can be accessed in this way only. But if we could know about mental qualities from the third-person perspective as representationalism holds, then quality inversion would be impossible. Since these intuitions depend on the claim in question, they cannot figure in good arguments for it (cf. Hilbert and Kalderon 2000). In other words, these intuitions covertly beg the question against representationalism by importing the central assumption of the traditional qualia theory.

Not only are there no good arguments for the traditional qualia theory, but there are also several reasons to think representationalism is preferable to it. Again, we are often aware of our mental qualities in consciousness. Any theory of mental qualities must respect and explain this fact. But common sense also holds that we often have access to others' mental qualities through the roles those qualities play in perceiving perceptible properties. For example, if someone picks

[97] The arguments in this chapter, and in particular in this section, are greatly influenced by Rosenthal's (2005) and (2010), as well as by discussions with him.

out a blue blueberry from a pile of red strawberries, then we have excellent evidence that the person's perception exhibits mental blue.

Moreover, though we are commonly right about what sorts of qualitative states we ourselves are in, we are, at least in some cases, open to correction about our own mental qualities. This is vividly illustrated by cases such as the phenomenon of dental fear, wherein dental patients report feeling pain upon drilling even though they lack the relevant nerves or have been anaesthetized. The following is a reasonable explanation of such a case: From the first-person perspective, it seems to these patients that they are in pain. In reality, these patients are not in pain. Rather, the patients are misconstruing their fear and sensations of vibrations as pain because they have a powerful expectation of pain (see, e.g., Rosenthal 2005, p. 209). If this is the right way to describe these sorts of cases, then there must be third-person ways of accessing others' mental qualities.

Any good theory of mental qualities must therefore explain not only our first-person conscious impressions of mental qualities, but also how such first-person conscious impressions integrate and interact with our third-person awareness of qualities. The traditional qualia theory cannot explain these things, but representationalism can do so.

Representationalism explains mental qualities in terms of their roles in representing perceptible properties. And the association of mental qualities with perceptible properties should not surprise us. Common sense has it that there is typically no difference between seeing perceptible red, for example, and having a sensation exhibiting mental red. Additionally, as noted in chapter 1, we describe mental qualities using the same predicates that we use to describe the perceptible properties of objects. We describe a sensation of red by using the very predicate 'red' or some related predicates such as 'reddish'. Indeed, how else could we describe mental

qualities, even from the first-person perspective? Our language for mental qualities is in that way parasitic on the language of perceptible properties.

The traditional qualia theory severs this tie between mental qualities and perceptible properties. In the attempt to accommodate phenomena such as quality inversion, traditional qualia theory holds that the connection between any given mental quality and any perceptible property is contingent at best. But if we only have first-person access to mental qualities independent of their connection to perceptible properties, how can we describe mental qualities? Proponents of the traditional qualia theory typically reply that, at bottom, we cannot describe them. One of the central features of mental qualities, qualophiles claim, is that they are ineffable and private. Though we may use predicates such as 'red' to describe a sensation, the sensation's connection to a perceptible red is optional. Hence Block famously claims, invoking Louis Armstrong, that if you have to ask about the nature of mental qualities, you just are not ever going to get to understand what they are. We have introspective or first-person access to them, and that is the end of the story.

There are additional reasons why one might think that mental qualities are not closely connected to perceptible properties. First, one might think that the fact that qualitative states can occur in the absence of perceptible properties (as in cases of hallucination) suggests that mental qualities can be characterized independently of perceptible properties. If there is no perceptible red in the vicinity when one tokens mental red, how could mental qualities be explained in terms of the perceptible properties? Moreover, one cannot perceive a perceptible property without first being in mental state with a given mental quality. This may suggest that, if anything, we should explain perceptible properties in terms of mental qualities, rather than the other way around.

Similarly, at present it seems that we cannot describe mental qualities in anything but the qualitative terms that we use to describe perceptible properties. By contrast, we can now describe perceptible properties in other terms such as in terms of their physical properties. This may give the impression that mental qualities and perceptible properties belong to two different explanatory domains. Perceptible red can be described not only in terms of that qualitative vocabulary, but it can also be described in terms of wavelengths of light. By contrast, mental red can be described as red, but it certainly does not appear to be, nor can we yet describe it in terms of, a physical property. We may hope to one day explain mental qualities in terms of brain states, but there does not seem to be much hope for that explanation at this time.

But all this shows is that, as Sellars says (adapting Aristotelian vocabulary), "visual impressions are prior in the order of being to concepts pertaining to physical color, whereas the latter are prior in the order of knowing to concepts pertaining to visual impressions" (1977, p. 177, emphasis his). Mental qualities arguably figure in the perceptual representation of perceptible properties, and in that fashion come before our concepts of perceptible properties in the order of being. Thus once we have a handle of on mental qualities' identities, we can understand them as occurring in the absence of perceptible properties. But if some version of representationalism is correct, we understand those mental qualities in terms of their role in the representation of perceptible properties—and so come after our concepts of perceptible properties in the order of explanation.

Proponents of the traditional qualia theory might reply that the conceivability of undetectable inversion nonetheless demonstrates that we have an understanding of mental qualities in consciousness independent of their connections to perceptible properties. But this intuition, as I've argued, is unfounded. Furthermore, our only models for how to understand

cases wherein mental qualities are in some way deviant from the norm depend on first understanding mental qualities in terms of their connections to mental qualities. Consider, for example, the ordinary phenomenon of color blindness. Our ability to detect and diagnose color blindness depends on differences between people's perceptual discriminatory abilities. It is because someone places red paint chips in the pile of green paint chips that we discover that the person is color blind. We understand what sort of deviant mental quality the person exhibits when he sees perceptible red in terms of what other perceptible colors he regards as similar to perceptible red.

Undetectable inversion goes well beyond this. The intuition of quality inversion requires that the differences between people lie in principle beyond what we can discover. It is thus unclear how we are to understand mental qualities, even from the first-person perspective. Theorists who attempt to rescue the traditional qualia theory simply assume that we have an independent way to characterize mental qualities via consciousness, but we do not possess such a way. The p-concepts approach might seem to offer that independent means by which to describe mental qualities in consciousness, but I've argued that these views are unworkable.

This is why representationalism, which aims to explain mental qualities in terms of their roles in the perceptual representation of perceptible properties, is attractive. Moreover, as I argued in chapters 1 and 2, there are many reasons to think that qualitative states do play this representational role. But, as I also argued in chapter 2, versions of representationalism which depend on simple tracking accounts of qualitative content are problematic because of their atomistic approach. One could conclude, as some representationalists have, that the intentionality of qualitative states is primitive, nonreductive, or otherwise resistant to informative explanation (e.g., Bourget 2010, Pautz 2010). But this reaction would be too hasty.

If we recall from chapter 2, perceptible colors are perceived to resemble and to differ from one another—and perceptible colors can even affect the way other colors are perceived. Moreover, as the cases of Swampman and Maxwell/Twin Maxwell illustrated, our perceptual discriminatory capacities are central to our ordinary ascriptions of qualitative content. What I think these facts show is that qualitative states are, at bottom, comparative or holistic phenomena. Since the kinds of tracking theories of content that standard versions of representationalism depend on are atomistic, versions of atomistic representationalism do not predict or explain these features of experience. Fortunately, there is a holistic theory of qualitative content available.

Recently, several philosophers have been converging on the idea that mental qualities are identified and individuated in a holistic way (e.g., Sellars 1963, Lewis 1972, Shoemaker 1975, Clark 1993, Rosenthal 2005).[98] In what follows, I'll sketch this view and propose a natural extension of it: Mental qualities are a species of nonconceptual content that are individuated holistically. The combination of this holistic theory of nonconceptual content with representationalism, I'll argue, avoids the problems that plague standard versions of representationalism and paves the way for a naturalistic theory of qualitative character.

§5.3 Towards a Holistic Theory of Qualitative Content

There are various ways to understand the proposal that qualitative states operate in a holistic way. Though the way that I will proceed to unpack this account fits comfortably with how

[98] In a recent paper, John Morrison (2013) defends a view that he calls 'anti-atomism about color representation', which may sound similar to this account. But, as Morrison acknowledges (p. xxx), his account differs substantially from views of the sort that I will discuss. On Morrison's view, perception involves occurrently representing relations between perceptible properties. As will become clear, the anti-atomist account that I'll develop does not involve these commitments.

several theorists have described it, I'll depart from their characterizations in various ways. Nevertheless, the view that I'll develop captures the basic shape of their views. In what follows, I will focus upon the perception of the perceptible colors, but my remarks arguably apply to perceptions and sensations involving other perceptible properties. The theory developed here may well apply not only to the perception of external proper sensibles such as perceptible color and smell, and to the external common sensibles such as shape and location, but also to the qualities of bodily sensations such as pains and itches.

§5.3.1 The Basics of Holism about Qualitative States

The foundation for this holistic account of qualitative states is the notion of a quality space (see, e.g., Shoemaker 1994, pp. 304-305).[99] Setting out to develop his theory of mental qualities, Shoemaker writes,

> A good place to start is with the truism that, on any view, creatures with the capacity for perceiving their environments will have "quality spaces," in Quine's sense of that term (see Quine, 1960 and 1969). For any given kind of such creatures, there will be physically different stimuli that are indistinguishable by creatures of that kind, in that they lack the capacity to respond differentially to such stimuli. And among the stimuli they can discriminate, some will be easier to discriminate than others. This imposes a similarity ordering on possible stimuli. The ordering will be relative to that kind of creature; the same stimuli may be ordered differently in the quality spaces of different kinds of perceivers (1994a, pp. 304-305).

It is plain that each sensory modality—such as vision, smell, or taste—enables a creature to distinguish among a family of perceptible properties. For example, we can see a range of perceptible colors. And it is well-established that we can chart what a creature can perceptually discriminate into what are known as quality spaces. A quality space is a map that reflects the ways in which the perceptible properties within a particular sensory modality are perceived to

[99] This is why Rosenthal calls his version of the view quality-space theory (e.g., Rosenthal 2010, p. 377).

relate to one another. Since perceptible red is more similar to perceptible orange than it is to perceptible green, perceptible red is closer to perceptible orange than it is to perceptible green in the quality space of the perceptible colors.

Importantly, as Quine observed, such innate quality spaces are arguably preconditions for explaining how creatures navigate the world and learn. These spaces reflect the similarities and differences between the properties that creatures can perceptually distinguish. If there were no such spaces, creatures could not perceptually distinguish anything and it would be at best unclear how creatures could interact with their environments. Quine writes,

> Without some such prior spacing of qualities, we could never acquire a habit; all stimuli would be equally alike and equally different. These spacings of qualities, on the part of men and other animals, can be explored and mapped in the laboratory by experiments in conditioning and extinction. Needed as they are for all learning, these distinctive spacings cannot themselves all be learned; some must be innate (1969, p. 123).[100]

At this point, it is important to note that to speak of the quality space of perceptible colors is somewhat misleading because there are many different quality spaces of perceptible colors, each of which is generated for different theoretical purposes. Color systems include, for example, the Munsell Color System, the Swedish Natural Color System, and the Oswald Color System (for discussion of these and other color-ordering systems, see, e.g., Hardin 1993, Cohen and Matthen 2010). These are all different ordering systems that organize the colors along distinct dimensions.

For example, the Swedish Natural Color System ("NCS") is based on color opponency theory, which holds that the light sensitive cells in the human eye process light in an antagonistic fashion—that is, the processing of light of one kind in a so-called opponent channel excludes the

[100] Though the theory is arguably committed to the view that some quality space(s) must be innate—that is, part of our inborn sensory endowment—it is neutral whether such spaces can be adjusted by learning or other developmental processes.

processing of light the other color in the channel. According to opponency theory, there are three opponent channels: Red versus green, blue versus yellow, and black versus white. The colors processed by the opponent channels are thus hypothesized to be so-called unique hues, hues that are not composites of any of the other unique hues. The NCS is thus generated by averaging out perceivers' judgments about the degree to which a given color exhibits a percentage of one or two of the unique hues, as well as its degree of blackness and saturation. As Rolf Kuehni reports, this process "results in a double cone system, formed by assembling constant-hue equilateral triangles of different hues along their common achromatic axis" (2010, p. 27).

The sort of quality space of perceptible colors upon which this holistic theory of qualitative states depends is the space of perceptible colors that best reflects a creature's capacity to discriminate among those perceptible colors. As Shoemaker notes above, however, any given creature will not be able to discriminate among every perceptible color. The human visual system is limited and there is a great diversity in people's perceptual capacities (for an interesting discussion of the various differences between the perceptual capacities of various groups, see Block 1999). But for any individual creature, there will be a quality space that best captures that creature's capacity to discriminate the perceptible colors.

Even though there is great variability regarding which perceptible colors individuals can discriminate, our commonsense conception of colors takes them to be mind-independent properties of objects. The fact that some people are colorblind and therefore cannot discriminate certain colors does not make us doubt that colors exist. We simply regard those people as being unable to see them. So, despite differences across individuals in terms of their perceptual discriminatory capacities, we can say that perceptible colors are whatever the human being who

can best discriminate can discriminate under standard lighting conditions. Thus, for any kind of creature, there is a space that averages over the kind to generate a space that captures the kind's discriminatory capacities. These are the sorts of spaces that I will refer to as the quality spaces of perceptible properties.[101]

There are thus distinct quality spaces for each family of perceptible properties distinguishable by means of every sensory modality. For example, the quality space of perceptible colors distinguished by vision is typically thought to be ordered along the three dimensions of hue, saturation, and brightness (see, e.g., Hardin 1993, pp. 25-26).

The central insight of this holistic view of qualitative states is that, in order for a creature to make such perceptual discriminations in the first place, the creature must be able to token a variety of psychological states that vary in ways that mirror the ways in which the perceptible properties relate to one another. As Shoemaker observes,

> Corresponding to the multidimensional similarity ordering on stimuli which constitutes a creature's quality space will be a similarity ordering on internal states of the creature produced by these stimuli, internal states that have the function of being perceptual representations of the external objects or states of affairs which, in the creature's environment, are the standard causes of the various sorts of stimuli. The similarities and differences amongst these states that enter into this similarity ordering are what I call "phenomenal" similarities and differences. These are functionally definable relationships. Roughly speaking, perceptual states have the phenomenal similarities and

[101] I note that it is a matter of debate how best to capture a creature's discriminatory capacities to thereby construct these spaces. At present, there are two ways to experimentally map a creature's discriminatory capacities: One can either test whether there are so-called just noticeable differences ("jnds") between two stimuli or whether the creature is simply able to distinguish them (for discussion of these two methods, see, e.g., Goodman 1977, pp. 197-200, 226-227; cf. Rosenthal 2005, p. 201, fn. 56). Ultimately, the view I sketch here is neutral about the best methodology for mapping a creature's perceptual discriminatory capacities—it is a question for perceptual psychology to answer. Similarly, it is for perceptual psychology to determine the dimensions of the quality spaces. As noted before, perceptible colors are typically thought to be ordered in terms the dimensions of hue, saturation, and brightness. But some color systems have other dimensions. For example, the NCS uses the dimensions of hue, chromaticity, and whiteness/blackness. Again, this view is neutral about what the dimensions might be, so long as the space that maps those dimensions best captures creatures' discriminatory abilities.

> differences they do in virtue of the roles they play in producing the behaviors that exhibit the quality space of the creature—the discriminations it makes, the kinds of conditioning it is capable of, the kinds of inductions it makes, and so forth (1994a, p. 305).

What are those psychological states in virtue of which creatures make perceptual discriminations? Since perception is qualitative and, as I've argued, there is no good reason to think that mental qualities play no functional role, it is natural to think that mental qualities are those psychological aspects of perceptions. Thus these theorists hypothesize that there are therefore spaces of mental qualities that are extrapolated from and match the corresponding spaces of perceptible properties. As Shoemaker describes it,

> Qualia I take to be the features of experiences in virtue of which they stand in these relations of phenomenal similarity and difference. They will stand to one another in relations of phenomenal similarity and difference in virtue of the roles they play in contributing to the phenomenal similarity or difference of the experiences that have them. The same experience may have many different qualia, and the overall similarity and difference relationships between two experiences will be a function of the similarity and difference relationships between the qualia instantiated in them. If, as I have said, the relations of phenomenal similarity and difference are functionally definable, then we should also be able to define in functional terms what the identity conditions of qualia are, and what it is to be a quale (1994a, pp. 305-306).

More recently, as Rosenthal puts it,

> [T]he quality space that reflects the perceptible properties a creature can discriminate by a particular sense modality will also determine the perceptual states that make such discriminations possible. And if mental qualities play a role in perceiving, it's natural to identify them with the properties in virtue of which perceptual states differ. The quality space that captures the similarities and differences among the perceptible properties a creature can discriminate will also describe the mental qualities that figure in such discrimination. And this gives us an account of mental qualities in terms of the quality space that describes a creature's ability to discriminate (2010, p. 377).

According to this holistic view of mental qualities, each mental quality can be identified and individuated in terms of its unique location with a space of mental qualities that matches its corresponding space of perceptible features. Just as the quality space of the perceptible colors reflects how perceptible red is more similar to perceptible orange than it is to perceptible green,

this holistic view holds that there is a matching quality space of mental colors in which mental red is more similar to mental orange than it is to mental green. Though individual mental qualities do not resemble perceptible properties, there is a resemblance between mental qualities and perceptible properties at the level of the families of properties. The family of mental colors resembles the family of perceptible colors.

On this view, each mental quality is thus identified and individuated in terms of its unique location within a space of mental qualities that matches a corresponding space of perceptible properties. Mental red is the property that it is because it occupies a unique location (relative to all other mental colors) within the space of mental colors that corresponds to the unique location (relative to all other perceptible colors) occupied by perceptible red within the space of perceptible colors. Put another way, mental qualities are identified and individuated in a holistic or relational way.

There are several versions of this holistic theory of mental qualities on the market.[102] At this point, it is worth mentioning an important difference between two prominent versions of it. As previously mentioned, the view arguably applies both to the proper and to the common sensibles. However, according to Austen Clark's feature-placing account, perceptions of common sensibles such as location do not involve mental qualities. Instead, Clark holds that

[102] Paul Churchland, for example, has also recently suggested a view according to which "our internal phenomenological color space is evidently a systematic homolog of the space of objective CA-ellipses. It is a reliable map of the global structure of that external feature space [of perceptible colors]" (2007a, p. 136). Though Churchland is notoriously skeptical about folk-psychological entities (e.g., Churchland 1981), he seems not to endorse eliminativism about mental qualities as he does for cognitive states. Churchland writes that "in one important area, a blanket eliminative materialism bids fair to be just plain wrong. The reason is simple: The portion of folk psychology concerned with the various sensations to which we are subject is in the process of finding a moderately smooth and highly illuminating reduction at the hands of unfolding neuroscience" (2007b, p. 166, emphasis his). But I will not explore Churchland's view in any further detail, however, because it is unclear what Churchland means by 'reduction' and therefore it is unclear whether he is giving a theory of mental qualities.

"[s]ensing proceeds by picking out place-times and characterizing qualities that appear at those place-times" (2000, p. 74, my emphasis). On Clark's view, these place-times are not determined by spatial mental qualities proprietary to each sensory modality, but by a more basic mechanism common to all sensory modalities.

According to Rosenthal's version of the view, by contrast, the perception of perceptible properties such as shape and location does involve mental qualities (see, e.g., Rosenthal 1999). On Rosenthal's view, the common sensibles such as shape are common insofar as the perceptible properties perceived by the various sensory modalities are the same. Both vision and touch can perceive the same perceptible shape. But Rosenthal holds that the mental qualities that figure in each sensory modality are proprietary to those modalities. The mental square quality in virtue of which one visually senses a square is a different property than the mental square quality in virtue of which one tactilely senses a square. Rosenthal thus gives a 'no' answer to what is known as Molyneux's question—the question of whether a man who was born blind but who had experienced tactile sensations of shapes would, upon being made sighted, be able to discriminate those shapes by vision (e.g., Locke 1700/1975, pp. 145-6; I, IX, 8). According to Rosenthal, any cross-modal calibration of this sort would have to be learned because the mental qualities via which one perceives shape tactilely are different from those that perceive shape visually (e.g., 2005, pp. 221-222).

There are reasons to prefer Rosenthal's version. Rosenthal argues that there is no reason to assume a common means of perceiving spatial properties insofar as different modalities perceive these perceptible properties in different ways. Vision, for instance, determines the boundaries of shapes in terms of the boundaries of colors, whereas touch determines shape boundaries in terms of the boundaries of texture. Moreover, there is evidence that cross-modal

calibration is learned. For example, Richard Held and colleagues (2011) found that patients who had recently received cataract surgery were capable of distinguishing pairs of legos separately by vision and by touch, but not capable of distinguishing a pair by vision and touch at the same time. These results are controversial (see, e.g., Schwenkler 2012), but I cannot settle these matters here and this difference between Rosenthal's and Clark's views are not important for present purposes.

§5.3.2 Holistic Representationalism Introduced

What does this holistic theory of mental qualities have to do with the contents of qualitative states? According to these views, mental qualities are aspects of perceptual states, aspects which track creatures' capacity to make perceptual discriminations. One could, nevertheless, maintain that creatures' discriminatory capacities generate quality spaces which in turn determine spaces of mental qualities, but that such mental qualities are not representational. On view, which would constitute a version of weak representationalism, each qualitative state exhibits perceptual representations in virtue of which it is able to make perceptual discriminations but also exhibits corresponding nonrepresentational qualia.

But this view is unmotivated for reasons discussed above. My argument for a version of strong representationalism is in that way abductive. Since there are no good reasons to adopt the traditional qualia view, it is reasonable instead to maintain that mental qualities are those mental properties in virtue of which we are able to discriminate perceptually among perceptible properties. That is, a version of strong representationalism according to which mental qualities are mental representations of perceptible properties is quite attractive. However, one could still maintain that creatures' discriminatory capacities generate quality spaces which in turn

determine spaces of mental qualities, but that such qualitative states' contents are nonetheless determined in an atomistic way. But such a view would face the above objections to atomistic accounts of qualitative content that I put forward in chapter 2.

The foregoing facts about qualitative states, however, suggest a more natural view of mental qualities—a theory that more neatly explains our perceptual discriminatory behavior than an atomistic account. Although proponents of this holistic account of mental qualities have not explicitly claimed that this account provides the semantics of qualitative representations, it is a reasonable extension of the view that the account provides a holistic theory of qualitative content.[103] On this view, a qualitative state represents a particular perceptible property just in case that state's mental quality (that is, its qualitative content) occupies a location within a space of mental qualities that corresponds to the location of that perceptible property within its space of perceptible properties. Playing on the name of holistic theories of conceptual content such as conceptual-role semantics ("CRS"), we can call this theory perceptual-role semantics ("PRS").

According to PRS, for example, the reason why mental red represents perceptible red, as opposed to perceptible green or some other perceptible color altogether, is because mental red occupies a relative location within the space of mental qualities that corresponds to the relative location occupied by perceptible red within the space of perceptible colors. That is, a qualitative state's content is determined in a holistic or relational way, in terms of its mental quality's

[103] Most proponents of this view bill it as a theory of 'qualia' or 'qualities', not as a theory of qualitative content. Nevertheless, these theorists do often make remarks that suggest that such qualities are a species of representational content (see, e.g., Rosenthal 2005, p. 208). For example, Rosenthal claims, "Qualitative states figure in perceiving things…. By reflecting the properties we perceive physical objects to have, qualitative states represent those perceived objects and properties" (2005, p. 208). Indeed, Clark acknowledges that his view can be read as a form of representationalism (e.g., 2004, pp. 571-572).

relative location in a space of mental qualities that matches the corresponding space of perceptible properties.

One could, however, read PRS simply as a theory of the contents of qualitative states—and remain neutral about whether those contents are (or determine) mental qualities. On this view, qualitative states sit in spaces of qualitative states that match corresponding spaces of perceptible properties, and so each qualitative state represents its corresponding perceptible property, but such representation has nothing to do with qualitative character. And PRS is quite attractive as a theory of the contents of qualitative states. For instance, PRS is fully consistent with many recent experimental findings about perception. As I have characterized it thus far, PRS explains the contents of qualitative states only in terms of their roles in perceptual discrimination, not in terms of our awareness of those contents in consciousness. PRS can thus accommodate the contents of qualitative states, even if those contents are exhibited by subliminal perceptions such as those that figure in patients with blindsight or normal individuals in masked priming studies.[104] Moreover, though PRS characterizes the contents of qualitative states functionally in terms of their roles in perception—and so makes no predictions about the neural realizers of those perceptual roles—PRS meshes nicely with the recent neuroscientific evidence demonstrating that there are distinct maps of neurons for each of the various modalities in the sensory areas of the brain (for smell, see, e.g., Howard et al 2009; for touch, see, e.g., Heed 2010; for taste, see, e.g., Chen et al 2011).

[104] This marks another important point of departure between, for example, Sellars's version of the theory and Rosenthal's. Rosenthal is explicit that his version is compatible with nonconscious perception (see, e.g., 2005, chapter 7). By contrast, Sellars seems to have held that all perceptions and sensations occur consciously. See, e.g., Sellars (1967, chapter 1), wherein he repeatedly refers to sense impressions as "states of consciousness" (at, e.g., p. 10).

But with these basic details of the view in hand, I'll now proceed to argue that PRS is the best theory of qualitative content—that is, the best basis for representationalism. I'll call a version of representationalism that depends on PRS holistic representationalism.[105] According to holistic representationalism, PRS provides the psychosemantics for the nonconceptual contents that are identical with (or determine) qualitative states' mental qualities.

§5.3.3 Some Benefits of Holistic Representationalism

If we recall, one of the main attractions of most standard versions of representationalism is that they depend on naturalistic theories of representation. Thus, by identifying mental qualities with particular kinds of representational contents, these versions of representationalism purport to offer naturalistic theories of qualitative character. As I noted in chapter 2, most standard versions of representationalism depend on atomistic accounts of content, accounts which putatively offer naturalistic explanations of representation. And, like those accounts, PRS provides a naturalistically respectable theory of the contents of qualitative states. What explains the contents of a creature's qualitative states—and thus the qualitative character of those states—are, at bottom, the creature's perceptual discriminatory capacities, which can be explained in purely physical terms.

Moreover, like standard versions of representationalism, holistic representationalism can arguably explain the qualities of all kinds of qualitative states. In other words, PRS can explain not only the qualities of sensations of external properties such as colors, but also arguably the

[105] As noted in chapter 1, since such holistic representation of perceptible properties can occur in the absence of consciousness, to give a complete account of the conscious qualitative character of experience, holistic representationalists must pair PRS with an adequate account of what makes the relevant qualitative contents conscious. Again, however, I remain neutral about which theory explains consciousness.

qualities of bodily sensations such as pains, tickles, and itches. As many representationalists have observed, if one attends to one's pain, for instance, one becomes aware of where the pain—that is, the bodily disturbance—seems to be located in one's body (e.g., Byrne 2001, pp. 227-229; Cutter and Tye 2011). PRS thus explains bodily sensations in the following way. Just as the space of mental colors matches the space of perceptible colors, so too does the space of what I will call mental pains (the qualities of the mental states of pain) correspond to the space of perceptible pains (the bodily disturbances).

Additionally, like standard versions of representationalism and unlike the versions of it discussed in chapter 3, holistic representationalism provides a way to understand the nature of mental qualities in a manner which fits comfortably both with our first and our third-person access to mental qualities. Since mental green is just that perceptual representation that enables a creature to perceptually discriminate perceptible green from other perceptible colors, we then recognize that, if a creature engages in behaviors that discriminate perceptible green from other perceptible colors, then we have excellent evidence that the creature's perceptions exhibit mental green. Likewise when we are aware of ourselves as being in states that enable those perceptual discriminations, we are aware of our own mental qualities. Holistic representationalism is thus to be preferred to the views discussed in chapter 4 because we need not posit peculiar p-concepts to account for special first-person access to mental qualities.

But there are many reasons to prefer holistic representationalism to atomistic versions of the view. First, holistic representationalism provides a tidy explanation of the commonsense evidence that chapter 2 discussed about the phenomenal structure of the perceptible colors. Again, atomistic representationalism possesses no resources on its own to explain this structure. By contrast, holistic representationalism explains the fact that perceptible colors are perceived to

resemble and differ from one another in characteristic ways in terms of the relative location of those perceptible colors within the quality space of perceptible colors that reflect one's space of mental qualities. The fact that perceptible red is perceived to be more similar to perceptible orange than it is to perceptible green is explained by perceptible red's closer proximity to perceptible orange than to perceptible green in the space of perceptible colors. In other words, just as holistic accounts of conceptual content predict and provide a straightforward explanation of the process of inference, whereas atomistic accounts of conceptual content do not, PRS gives a straightforward explanation of why perceptible properties perceptually resemble one another. PRS gets those resemblance relations for free.

Moreover, recall that atomistic representationalism runs into difficulties because it must invoke notions such as optimal conditions to explain cases of misrepresentation. The theory cannot, for example, explain how Swampman is able to perceive perceptible colors. Because Swampman has no evolutionary history, he cannot stand in any tracking relation (unpacked in terms of biological function) to the perceptible colors. PRS, by contrast, does not invoke notions such as optimal conditions to explain cases of misrepresentation. According to PRS, a qualitative state misrepresents some perceptible color simply if one perceptually discriminates that perceptible color as occupying the wrong location in one's quality space of perceptible colors (determined by one's dispositions to make the relevant perceptual discriminations). If I perceptually discriminate something perceptible red as occupying the location of perceptible blue in my quality space of perceptible colors, then I misperceive that perceptible red thing as perceptible blue. One's capacity to perceptually represent perceptible colors thus does not depend on one's evolutionary history—rather, it depends on one's occurrent capacity to perceptually distinguish the perceptible colors. So if Swampman has the perceptual ability to

discriminate a quality space of perceptible colors, then PRS holds that he is able to perceive those perceptible colors.

Likewise, holistic representationalism delivers the correct result in Pautz's case of Maxwell and Twin Maxwell. Since both Maxwell and Twin Maxwell track the same property, atomistic representationalism must maintain that they be in states that exhibit the same mental qualities. This is implausible, however, because Maxwell and Twin Maxwell exhibit such different discriminatory behaviors. But according to holistic representationalism, one's discriminatory capacities constitutively determine what one represents and thus determine one's qualitative character. Since Twin Maxwell locates perceptible-orange objects in the perceptible red region of his quality space, holistic representationalism delivers the result that, even if he is tracking perceptible orange in optimal conditions, he is misrepresenting it as perceptible red and thus perceiving perceptible red.

In addition, holistic representationalism both predicts and provides the beginnings of an explanation of a variety of experimental evidence, such as the evidence of color-context effects. Recall, PRS holds that a mental quality's content is determined by that quality's relative location within a space of mental qualities. So each qualitative state includes relative information regarding all of the other mental qualities that one can exhibit. A reason to expect perceptible color-context effects, then, is that this relative information built into each qualitative state causes mental qualities, as it were, to bias other mental qualities. For example, a perception of a dark background causes one's perception of a blue stimulus to appear lighter because the biasing effect of the perceptible colors of the background. This hypothesis is merely a sketch, but the core idea is that color-context effects can be explained by the interactions of the complex relational information embedded within each mental quality.

Consider, by analogy, the way that thoughts can bias one another. It is well-established experimentally that having thoughts about dogs biases one to have thoughts about things conceptually related to dogs, such as cats. This is the phenomenon known as semantic priming (see, e.g., McNamara 2005). Holistic theories of conceptual content such as CRS provide a nice explanation of such biasing effects: Because of the inferential connections between thoughts about dogs and subsequent thoughts (connections which characterize those thoughts' contents), one's having a thought about dogs disposes one to form those subsequent thoughts. Likewise, because of the relative information regarding other mental qualities encoded in each mental quality, a particular quality can bias other mental qualities, resulting in context effects.

Moreover, unlike versions of conceptualist representationalism discussed in chapter 2, holistic representationalism permits ready explanation of the distinctions between thought, perception, and sensation. As we recall from chapter 2, one of the main reasons for positing that qualitative content is nonconceptual is that it enables us to draw distinctions between these varieties of mental states. In the following section, I'll argue that mental qualities are a species of nonconceptual content. On this view, sensations exhibit mental qualities only, thoughts exhibit conceptual properties only, and perceptions (and arguably other states such as emotions) are hybrid states exhibiting both qualitative and conceptual properties.

§5.4 The Nature of Qualitative Representation

PRS is an account of the semantics of the nonconceptual content of qualitative states. In this section, I will spell out how this nonconceptual form of qualitative representation—that is, representation via mental qualities—differs from conceptual representation, which is representation via conceptual states.

As I've argued, there are many reasons to think that the contents of qualitative states are nonconceptual. As noted in chapters 1 and 2, perhaps the best reason to think that qualitative states exhibit nonconceptual content is that such an account explains the presence of what I've called perceptual thoughts—that is, thoughts about what one is currently sensing or perceiving. It is likely that there are relatively simplistic creatures who are capable of being in qualitative states, but that are incapable of having thoughts. However, for those creatures that are capable of thought, an excellent explanation of the presence of perceptual thoughts is the fact that creatures sense or perceive what those thoughts are about. Indeed, it is typically assumed that if a creature exhibits a qualitative state, then it is disposed to form a perceptual thought about the perceptible properties that correspond to that state. If I have a sensation of red that exhibits mental red, for example, I am inclined to think that there is perceptible red. A natural explanation of the existence of perceptual thoughts and this close connection between sensation and thought is that both mental qualities and perceptual thoughts represent perceptible properties, but in different ways. Mental red grounds the disposition to have the thought that there is perceptible red because both the mental quality and the perceptual thought represent perceptible red.

But how are we to understand the nonconceptual representational character of mental qualities? Since qualitative content is proposed to explain a species of conceptual representations (perceptual thoughts), it follows naturally that we can understand mental qualities' representational characters on the model of the representational properties of those conceptual representations. In other words, I offer here something analogous to Sellars's claims that, when it comes to perceptible properties, mental qualities come first in the order being, though perceptible properties come first in the order of explanation. In a similar fashion, since representational mental qualities cause perceptual thoughts, mental qualities come before

perceptual thoughts in the order of being. But I argue that we should understand the representational character of mental qualities in terms of our understanding of the representational properties of perceptual thoughts—and in that way perceptual thoughts come before the representational character of mental qualities in the order of explanation.

Insofar as perceptual thoughts are cognitive states, a perceptual thought exhibits a mental attitude towards a conceptual content. But what kind of content and what kind of attitude? A perceptual thought concerns a perceptible property that one currently perceives, so it does not merely represent that perceptible property, but represents it as present. Thus the content of a perceptual thought is, roughly, *that there is some perceptible property present*. One can, of course, exhibit several different mental attitudes towards this representational content. One can, for example, hope that there is some perceptible property present. But the hope that there is a perceptible property present does not commit one to the presence of that perceptible property, whereas the perceptual thought that there is that perceptible property present does. So the attitude of a perceptual thought is not, for example, hoping, doubting, or wondering, but rather the assertoric attitude. The perceptual thought that there is perceptible red present, for example, asserts the content *that there is perceptible red present*.

Because a perceptual thought's content exhibits this sentence-sized character, it represents something like a state of affairs—namely, the state of affairs that there is some perceptible property present.[106] Because perceptual thoughts represent things like states of affairs, their contents are thereby truth evaluable. If the state of affairs that perceptible red is present obtains, then the content of one's perceptual thought is true, and it is false otherwise.

[106] I will not take a stand on precisely what cognitive states represent (e.g., facts, states of affairs, etc.). My point is only that, whatever they represent, they do not represent properties alone, as I'll argue mental qualities do.

Moreover, in representing these states of affairs, perceptual thoughts thereby engender an awareness of those states of affairs. Of course, it may be that one cannot be aware of anything if that thing does not exist. In other words, the awareness relation is arguably factive. But when the state of affairs that a perceptual thought represents obtains, the perceptual thought makes one aware of that state of affairs.

Like perceptual thoughts, mental qualities represent perceptible properties. But unlike perceptual thoughts, mental qualities do not represent things like states of affairs. Rather, mental qualities represent perceptible properties only. This is why we describe mental qualities with term-sized predicates such as ‘red’ or ‘reddish’ instead of with sentential descriptions. Mental red represents perceptible red only. But mental qualities do not represent perceptible properties in just any way. Like a perceptual thought, a mental quality represents its corresponding perceptible property in a way that commits one to the presence of that perceptible property. For example, if one tokens mental red, one is committed to the presence of perceptible red. Importantly, however, mental qualities do not represent perceptible properties as being present. Representing a perceptible property as being present is equivalent to representing the state of affairs that the perceptible property is present. Again, mental qualities represent properties only.

So how do mental qualities commit us to the presence of perceptible properties without representing those perceptible properties as being present? Because mental qualities are not conceptual, they do not exhibit anything akin to mental attitudes. One cannot, for example, assert—let alone hope, doubt, or wonder—some mental quality. Though mental qualities do not exhibit mental attitudes, we can understand the way in which mental qualities commit us to the presence of perceptible properties as a qualitative counterpart to the assertoric attitude of perceptual thoughts. This assertoric aspect of mental qualities is not, strictly speaking, a mental

attitude. It cannot be taken towards other representational contents, nor are there distinct varieties of it. All mental qualities have this assertoric character only. The way in which mental qualities commit us to the presence of their corresponding perceptible properties is built into the representational character of those qualities.

Since mental qualities represent properties only, they cannot be true or false in the way that the conceptual contents of perceptual thoughts can be. Rather, mental qualities exhibit something akin to truth conditions, which I've called accuracy conditions (as I noted in chapter 1, truth conditions and accuracy conditions are both kinds of veridicality conditions). For example, if one exhibits mental red, one's mental quality is accurate just in case perceptible red is present and inaccurate otherwise. And just as perceptual thoughts engender an awareness of states of affairs, mental qualities engender a different, but related, kind of awareness of perceptible properties. When the perceptible property that a mental quality represents obtains, the mental quality makes one aware of that perceptible property. To borrow a bit of terminology from Dretske, perceptual thoughts engender what he calls fact awareness—that is, an awareness of things like states of affairs—whereas mental qualities engender what he calls property awareness, an awareness of properties only (1999, p. 106). Dretske's distinction mirrors the commonsense distinction between being aware that something is the case (fact-awareness) and being aware of something (property or object awareness).

One might object that both mental qualities and perceptual thoughts render us aware of perceptible properties. After all, common sense has it that there are two ways to become aware of something. First, one can become aware of something by sensing it. If I sense perceptible red, I am thereby aware of perceptible red. And, as I've been arguing, the sensation of perceptible red makes one aware of perceptible red because the sensation exhibits mental red,

which represents perceptible red. But one can also become aware of something by having a thought about that thing as being present. If I have the thought that there is a red tomato in front of me, I am arguably aware of that red tomato. Likewise, one might argue that perceptual thoughts, in making us aware of the states of affairs that there are perceptible properties present, also make us aware of those perceptible properties.

Perceptual thoughts nonetheless represent perceptible properties in a different way than do mental qualities. Employing Dretske's terminology, though it is true that some fact-awareness does constitute a kind of property-awareness (e.g., one's fact-awareness of the states of affairs that there is perceptible red present involves a property-awareness of perceptible red), the sort of property-awareness that mental qualities engender does not constitute a fact-awareness. In other words, while some awareness that something is the case is also an awareness of something, not all awareness of something is awareness that something is the case.

There are several additional differences between qualitative representation and conceptual representation that are worth noting here. For example, one important feature of conceptual contents is that they are amodal in at least two ways. First, one can have a thought about anything—including thoughts about things that are not distinctively sensory. We do not describe the thought that 119 squared equals 14,161 as being in any way related to or derived from any sensory modality. Second, the same conceptual content can figure in any kind of cognitive state such as a belief, a hope, or a perception. If cognitive states are composed of concepts, then the same concept RED figures in both the thought that there is something red present and the perception that there is something present.

Mental qualities, by contrast, are modality specific. Mental red, for instance, is distinctively tied to vision and thus can only figure in visual sensations and visual perceptions of

perceptible red. Mental red cannot figure in a thought about perceptible red, whereas the concept RED can. And this of course is consistent with how mental qualities appear from the first-person perspective. Introspection seems to reveal that mental qualities stand alone and do not seem to be able to fit together with other qualities or concepts. Even though mental qualities have a term-sized character as concepts do, mental qualities cannot combine with one another to form sentence-sized units. Mental red cannot combine with mental green, though it can be mentally located next to mental green.

The fact that both mental qualities and concepts exhibit term-sized character may explain why so many philosophers have pursued the Aristotelian/Humean project of building representational content out of mental qualities, a project which I briefly discussed in chapter 1. But because mental qualities cannot combine with one another or with concepts to generate conceptual contents, whereas concepts do, there is at least one reason to doubt this Aristotelian/Humean position. Another reason to doubt this Aristotelian/Humean view is that it is an important feature of cognitive states that they can represent things as existing or as not existing (cf. Tye 1995, p. 108). The belief that the dog is brown is true just in case the state of affairs that the dog is brown obtains and false otherwise. Just as the belief that the dog is brown represents the dog as existing, the belief that there are no unicorns represents unicorns as not existing. Mental qualities, however, cannot represent things as not existing. Again, mental qualities represent perceptible properties only. Though mental red can inaccurately represent perceptible red even if no perceptible red is present, mental red cannot represent perceptible red as not present. This also renders the picture on which content is built out of qualities questionable (cf. Rosenthal 2012, p. 24, fn. 3). If mental qualities cannot represent things as not existing, it is unclear how they could compose contents that do represent things as not existing.

Because of the sentential character of conceptual contents, cognitive states can stand in inferential relations to one another in respect of their contents. One can infer from perceptual thought that there is red to the belief that something is colored, and so on. This is why many identify concepts as the inferentially relevant constituents of cognitive states (e.g., Crane 1992, p. 147). By contrast, qualitative states cannot stand in inferential relations in respect of their mental qualities. One cannot infer anything from mental red. Mental red does not imply, for example, some mental equivalent to nongreen. One can, of course, infer something from the thought that one's state exhibits mental red, and one's thought can be true or false, but that is because the thought has the conceptual content *that one has a reddish sensation*. If one never has sensations that exhibit both mental red and mental green at the same mental location, it is because of the nature of those mental colors, not because of some inferential relation that holds between them.

But perhaps the biggest difference between conceptual and qualitative representations is that different theories explain why these represent what they do. As noted in chapter 1, there are a variety of theories proposed to explain why conceptual states exhibit the contents that they do, including causal-covariational theories (e.g., Stampe 1977, Dretske 1981), CRS (e.g., Harman 1987), and asymmetric-dependence theory (e.g., Fodor 1990). I will not attempt to decide between these theories. But whichever theory of conceptual content proves to be correct, these theories do not explain why mental qualities represent what they do. Clearly atomistic theories of conceptual content, which attempt to explain why cognitive states exhibit the content that they do independent of the connections between those contents, cannot explain qualitative representation. Qualitative representations, I've argued, are individuated holistically. But even holistic accounts of conceptual content such as CRS, which explains cognitive states' contents in

terms of their inferential relations to one another, are quite unlike PRS. According to PRS, by contrast, mental qualities do not stand in inferential relations to one another. Rather, mental qualities stand in relations of resemblance and difference, determined by the resemblance-and-difference relations that characterize a creature's corresponding space of perceptible properties. Conceptual contents are not determined in this way.[107]

Some have attempted to explain conceptual representations in terms of their relations to one another in a way that does resemble PRS in some respects. For example, Peter Gärdenfors argues that we can understand various concepts in terms of their locations in what he calls conceptual spaces. Conceptual spaces are defined by their dimensions, and the dimensions of spaces are purportedly either innate or given by theories, social conventions, and so on (e.g., Gärdenfors 2004, pp. 15-18). For example, speaking of Newtonian mechanics, Gärdenfors writes,

> The quality dimensions of this theory are ordinary space (3-D Euclidean), time (isomorphic to the real numbers), mass (isomorphic to the non-negative real numbers), and force (3-D Euclidean space). Once a particle has been assigned a value for these eight dimensions, it is fully described as far as Newtonian mechanics is concerned. In this theory, an object is thus represented as a point in an 8-dimensional space (2004, p. 18).

This is an intriguing proposal. But this proposal differs from PRS because the quality dimensions of Gärdenfors' proposed conceptual spaces are generated by theories of mathematics and physics (as examples), not by the matching spaces of similarities and differences among corresponding properties or states of affairs. Moreover, this proposal cannot work in the way

[107] This line of reasoning provides a basis for a new argument for the existence of nonconceptual content. Whatever account of conceptual content is true (whether holistic or atomistic), that account does not hold of qualitative content—and thus qualitative content is nonconceptual (this argument was brought to my attention in conversation by Prinz).

that PRS does because nonperceptible properties do not fall into spaces that differ along particular dimensions. As Rosenthal observes,

> Dogs bear no suitable relations of similarity and difference to everything else we can form a concept of; so we can't appeal to such relations to fix the concept of a dog.... Concepts often apply to things whose natures fall into no such well-defined families (2005, p. 208).

Unlike perceptible properties, nonperceptible properties such as the property of being a dog do not sit in circumscribed families of similarity-and-difference relations. The spaces of perceptible properties that differentiate a creature's spaces of mental qualities are only quasi-holistic, insofar as they are circumscribed families of properties. A creature's quality space of perceptible colors reflects relations holding between the perceptible colors only. Concepts or thoughts, by contrast, do not fall into families that correspond to such circumscribed orderings of properties or states of affairs. Any given concept stands in some inferential relation to any other concept. However it is that conceptual representation operates, such an operation differs from the way in which qualitative representations function.

§5.5 Putting Holistic Representationalism to Work

If mental qualities are representational in the way that I propose, then we have the beginnings of explanations of many phenomena.

§5.5.1 Explaining How Things Seem

For example, in chapter 1 I briefly noted that whenever one's qualitative states change, it appears that the way things seem to one change as well. For example, if I shift from being in a state that exhibits mental red to a state that exhibits mental green, then at first things seem to me to be perceptible red and next things seem to me to be perceptible green. If how things seem to one is

a matter of how one is representing things, then this argument appears to establish a form of strong representationalism. Following Lycan (2006), we can call this argument The Argument from Seeming for representationalism (see, e.g., Byrne 2001, p. 211; Thau 2002, pp. 32-33; Schellenberg 2011, pp. 719-720; for discussion, see, e.g., Lycan 2006, §3.4 or Crane 2009b, pp. 484-486).

But as discussed in connection with Shoemaker's and Block's views in chapter 3, it is doubtless that expressions such as 'looks', 'appears', 'seems', and their cognates are ambiguous in a variety of ways. The foregoing analysis of the representational character of mental qualities can help explain at least one use of expressions such as 'looks' or 'appears'. When one beholds a white piece of paper partially bathed in red light, one might say that the paper looks or appears to be partially red, even though one can see that it is white. If, however, one viewed the white paper bathed in red light, but did not know that the lighting conditions were abnormal and instead supposed that the conditions were normal (that is, one supposed that the paper were bathed in white light), one might erroneously think or say that the paper is red. What could account for the fact that one has such a false thought about the color of the paper? According to the present account, the explanation is that the red light causes one to token mental red, which represents perceptible red. Because one does not know about the strange lighting conditions, the presence of mental red causes one to form the false thought that it is the paper that exhibits the perceptible red.

Now, suppose that one did know or had reasons to suspect that the lighting conditions were abnormal. In that case, one would not be inclined to think or say that there is a red piece of paper. Instead, one would be inclined to think or say that the paper seems or appears to be red. Of course, one might have been inclined to think or say the same thing when one did not know or

suspect the lighting conditions were nonstandard as well. As noted in chapter 2, it is a folk psychological truisim (that I called the perceptible-color truism) that if something is, for example, perceptible red, then it looks under standard lighting conditions to normal observers to be perceptible red. When someone says that an object looks or appears or seems to be some perceptible color, one typically does so to signal that one is inclined to take it to be that perceptible color, but one is suppressing that inclination because one is unsure what the lighting conditions are, and so wishes to withhold judgment about the object's perceptible color. After all, if the lighting conditions were nonstandard, it might be the case the object's color is not as it appears. When one does not have any reason to suspect that the lighting conditions are abnormal, one simply can say that the object's color is as it appears.

A natural explanation of why one has the inclinations that one does is that when one says that something seems a way, one exhibits the same mental quality that would accurately represent that perceptible color were the lighting conditions normal. Since mental red is the mental property that represents perceptible red, mental red disposes one to form the thought or say that red is present. If one has no reason to suspect the lighting condition to be unusual, one simply attributes that perceptible color to the relevant object. When one says that something looks red, one distrusts one's inclination to attribute that perceptible color to that object.

Moreover, it is reasonable to think that this is how we learn to use expressions such as 'looks' or 'appears' in the first place. A natural hypothesis is that young children start out regarding all objects to be the perceptible colors they appear to be. So a child would simply take a red object bathed in blue light to be purple. This is because mental colors are the natural reaction to perceived colors, which dispose people to render judgments that those perceptible colors are present. Only after learning about the vagaries of lighting conditions and other similar

influences do children come to find that their initial judgments about the perceptible colors of objects may be false, and so learn to hedge those judgments by using expressions such as 'looks' instead of 'is' in situations of doubt. In this way, 'looks'-talk is parasitic on 'is'-talk; we come to acquire the former by first becoming fluent with the latter (cf. Sellars 1963, pp. 149-156).

Recall, by contrast, Shoemaker's suggestion that something's looking red involves its exhibiting a distinct property than its perceptible color, which he calls appearance red. That is, Shoemaker rejects the perceptible-color truism. But I argued in chapter 3 that Shoemaker's account renders it impossible to give any informative account of the perceptible colors. The same difficulty does not, however, plague the present account. On this view, something's looking red to one consists in one's exhibiting mental red, and thereby representing perceptible red to be present, though one withholds judgment about whether that thing actually is perceptible red. So something's looking red does not involve attributing properties distinct from the perceptible colors to the object. The only kinds of colors are the mental colors and the perceptible colors—there are no appearance colors. When one says that something looks red, what this amounts to is not that one perceives the object to exhibit some appearance property other than its perceptible color, but rather that one exhibits mental red. Perception does not represent appearance properties.

In chapter 3, I noted that there is debate regarding whether there is a representational use of 'seems'. Shoemaker, for example, argues that uses of 'seems' involve representing things, whereas Block urges that at least one use of it tracks one's mental qualities and not its representational properties. And some have argued that, if there is such a qualitative use of 'seems', then the Argument from Seeming fails to establish representationalism (see, e.g., Crane 2009b, p. 484). If there is such a phenomenal sense of 'seems', then the argument trivially

asserts that there can be no change in the qualitative character of one's experience without a change in one's qualia, a claim which hardly establishes representationalism.

On the present account, however, at least one use of 'seems' tracks both one's mental qualities and one's representational properties. To say that something seems red is to say that one is representing perceptible red via mental red. Block and Shoemaker are thus both partially correct. Shoemaker is correct to assume that things' seeming a way is a matter of representation and Block is correct to assume that things' seeming a way is a matter of what mental qualities one is tokening. But they are both wrong to assume that a claim about one's mental qualities is not also a claim about what one is representing. This is why the Argument from Seeming for representationalism succeeds.

§5.5.2 Grounding and Justification

As discussed in chapter 2, many have argued that we must attribute content to sensation and to perception in order to explain how sensations and perceptions ground or justify perceptual thoughts (e.g., McDowell 1994, Pautz 2010a). In that chapter, I argued that this argument rests on the assumption that the only properties that can ground or justify perceptual thoughts are conceptual contents, and that this assumption is questionable. And the reason why mental qualities can perform these functions should be plain. As I've been arguing, mental qualities' representational characters are posited in the first place to explain the presence of perceptual thoughts.[108]

[108] Likewise, I noted in chapter 2 that many posit nonconceptual content to explain how nonhuman animals and infants can represent the world without possessing conceptual capacities and to explain how we acquire concepts in the first place. Similar arguments for how mental qualities can play these roles can be given as well.

Though one might grant that mental qualities can ground perceptual thoughts, one might nonetheless deny that mental qualities can justify those thoughts. After all, the psychological relationship that I argue holds between mental qualities and perceptual thoughts is quite unlike the relationship that holds between one conceptual state and another. Conceptual states stand in inferential relations to one another, whereas mental qualities cannot stand in inferential relations to one another or to conceptual states. However, it is arguable that qualitative states are related to perceptual thoughts in a psychologically robust way that is akin to inference. While one does not, strictly speaking, infer from mental red the perceptual thought that there is perceptible red, the quality does typically cause the thought. Because of this regular causal connection between mental qualities and perceptual thoughts, mental qualities ground and justify their corresponding perceptual thoughts.

I acknowledged in chapter 2 that one could dig in one's heels to insist that justificatory relations hold only between such conceptual states as can enter into logical relations. But I also argued that it is not clear why one should insist upon this. And even if it were the case that only a conceptual state can justify another conceptual state, we can nonetheless recognize that there is a genuine epistemic relationship that holds between mental qualities and perceptual thoughts. The idea that there is an epistemic relationship between nonconceptual states and conceptual thought is not new. Edmund Husserl, for example, called this sort of relationship epistemic fulfillment (e.g., 1970, Investigation 6, §8). Husserl claims that a perception can fulfill a thought, even though a perception and a thought exhibit different structures, because they pertain to the same target.[109]

[109] On Husserl's view, perceptions are arguably wholly nonconceptual, whereas on my usage of 'perception', perceptions always involve conceptual content. But the details of Husserl's account here are not necessary, though the basis idea of his proposal is instructive. For an

I propose, then, to appropriate Husserl's notion of fulfillment to describe the epistemic relation between mental qualities and perceptual thoughts. Since a mental quality and its corresponding perceptual thought represent the same perceptible property, we can say that the mental quality fulfills the perceptual thought. Consistent with the Quinean sort of naturalized approach to epistemology that I discussed in chapter 2, fulfillment is the only kind of justification that epistemology should require from sensation.[110]

In the developmental course of an individual, mental qualities may not immediately begin to cause the relevant perceptual thoughts. But as one repeatedly encounters the relevant perceptible properties, and thereby tokens the relevant mental qualities, one's dispositions to react to one's environment will become ever more varied and complex. As this process reoccurs and is refined, one's mental qualities give rise to the relevant perceptual thoughts. Eventually, the transition between the mental quality and the corresponding perceptual thought becomes quick and automatic insofar as the perceptual thought practically accompanies the sensation. And this is what it is to perceive something. Perception is just a perceptual thought accompanied by sensation. At this point, I will remain neutral regarding whether perception is a hybrid state

---

overview of Husserl's proposal, see, e.g., Hopp 2011, chapter 7. The present discussion about how mental qualities justify perceptual thoughts is influenced by Hopp's discussion of fulfillment.

[110] In holding that perceptual thoughts are justified (or fulfilled) by mental qualities, I am not endorsing indirect realism of the sort championed by Bertrand Russell (1912/1997) or A. J. Ayer (1956). According to that sort of view, what one strictly perceives is something mental—a so-called sense-datum—and one makes an inference that the nonmental object that causes that sense-datum is present. This view has long been thought to be problematic for a variety of reasons (see, e.g., Austin 1962a). On my proposed view, sensations make one aware of nonmental perceptible properties. These sensations then cause one to form perceptual thoughts about the very same nonmental perceptible properties that the mental qualities represent. In order to perceive perceptible properties, one need not perceive anything mental or draw inferences from perceptions of anything mental. As Sellars puts it, "The direct perception of physical objects is mediated by the occurrence of sense-impressions, which… are, in themselves, thoroughly non-cognitive…. Sense impressions do not mediate in virtue of being known" (1963, pp. 90-91).

that encompasses both a sensation and a perceptual thought, or whether these are two distinct states that are causally related in such a way that they coincide temporally.[111]

§5.5.3 A Two-Component View of Perception

Perception thus involves two representational components: The sensory component of perception qualitatively represents perceptible properties, and the thought component of perception conceptually represents whatever is perceived (which may include but is arguably not limited to perceptible properties[112]). In that way, perception is, as some have argued, layered (e.g., Lycan 1996b, Tye 1996).[113]

Before cataloguing the virtues of this two-component view of perception, I'll first note what might seem to be an objection to it. It does appear that perception represents coherent three-dimensional objects and not a jumbled assortment of perceptible properties such as colors

[111] It is thus compatible with my proposal that one can have a sensation of a perceptible property without the corresponding perceptual thought about that property, and thereby not perceive that property. Such cases are presumably rare, but there do appear to be examples of them. Suppose, for example, that one has a brief hallucination of something red when awakening from sleep. Because of the confusion that often accompanies awakening, one may not have the time or wherewithal to form thoughts about anything. All that happens is that there is a flash of red, as it were, across one's visual field. Since there is no (or since there need not be) a red object around, one does not assert that anything is the case; one does not assert, for example, that redness is present. A natural explanation of this kind of case is that one senses, but does not perceive, red (for a description of a similar case, see, e.g., Bengson et al 2011, pp. 170-171).

[112] There is a question regarding what sorts of things perception does conceptually represent. I will address this question in §7.1 in connection with the so-called debate over the contents of perception.

[113] This account differs from most accounts of the layers of perception in several ways, however. For example, Lycan argues that perception is layered insofar as we perceive certain things in virtue of perceiving other things. Lycan writes, "We do visually represent the trees, and as being of the same size etc., but we do that by representing colored shapes and relations between them" (1996b, p. 95). But it is not clear how Lycan understands this relation between representations. On the view defended here, the only way in which one conceptually represents that there is red in virtue of qualitatively representing red is that the qualitative representation causes the conceptual representation.

and shapes. But if, as I've been arguing, mental qualities involved in perception represent only individual properties, why does perception appear coherent in this way? This is a question closely related to the so-called binding problem (e.g., Treisman 1998), which concerns whether the features present in perception are unified and, if so, in what way. According to the view developed here, the implication for the binding problem is that whatever binding may take place does not take place at the level of qualitative representation, but at the level of conceptual representation. It is true that certain mental qualities invariably accompany other mental qualities. So, for example, mental shapes are never present without the presence of mental colors. But this is not because mental colors are somehow always psychologically bound to mental shapes; their relation is simply that they happen to be coinstantiated.

Mental qualities such as mental shapes and mental colors may seem to be bound, but this appearance can be explained otherwise. First, mental shapes and mental colors are invariably coinstantiated at the same mental locations, which may give the impression that they are bound. That mental shapes and colors are invariably coinstantiated can be explained by the fact that it is in the nature of perceptible shapes that they are never present in the absence of perceptible colors. Perceptible shapes and colors are collocated in the world, and thus their psychological counterparts are never instantiated independently of one another (cf. Rosenthal 2005, p. 200). Second, in perception, one typically conceptually represents the perceptible properties that one also qualitatively represents as being properties of a unified object. For example, if I perceive that there is a dog, I qualitatively represent all of the disparate perceptible properties that I perceive, such as the perceptible shape and color of the dog. But I also conceptually represent those perceptible properties as being properties of the dog, which may give the impression that

the mental colors and shapes are somehow psychologically bound at the level of qualitative representation, when they are only integrated at the level of conceptual content.

This two-component view of perception does have many explanatory virtues. For example, in chapter 2 I observed that perception involves aspects that are inferentially isolated (e.g., the discussion of Crane and the waterfall illusion). The two-component view nicely explains this feature of perception. Because mental qualities cannot figure in inferences, one's beliefs typically do not alter what one qualitatively represents in the same way that one's beliefs can alter what one conceptually represents.[114]

The two-component view of perception also explains the phenomena of perceptual constancies. It is a well-known feature of perception that people veridically perceive objects as being uniform in, for example, color, even though those objects may vary in illumination and thus vary in the physical spectra of light striking the retina (for discussion, see, e.g., Palmer 1999, pp. 312-313; Burge 2010, chapter 9). For example, a white piece of paper known to be partially bathed in red light is typically perceived to be wholly white, even though some red light does strike the retina. A natural explanation of this sort of phenomena is that, while one conceptually represents the paper to be uniformly perceptible white—that is, one has the perceptual thought that the paper is uniformly perceptible white—one qualitatively represents several distinct perceptible properties, including both perceptible white and perceptible red.

Even though there are two components of perception representing different perceptible properties, the perception of constancies in this way need not be falsidical. In the case of the white paper partially bathed in red light, for example, some red light is striking one's retina. So

[114] I say 'typically' because I do not rule out the possibility that beliefs (or other conceptual states) can interact causally—that is, in a noninferential way—with mental qualities to affect how one qualitatively represents things to be.

one's qualitative representations are not inaccurate because some perceptible red is present. Thus one's perception is wholly veridical: The perceptual-thought component represents the paper as being perceptible white (and its content is therefore true), whereas the sensory component represents the presence of both perceptible white and perceptible red (and its mental qualities are therefore accurate).

§5.6 Objections

This section considers a variety of objections to holistic version of representationalism. Defusing these objections will not only show how they do not threaten the view, but also help to elucidate it.[115]

§5.6.1 Mental Qualities Appear Intrinsic

I'll first consider a fairly basic objection to holistic representationalism and the holistic theory of content, PRS, on which it depends. Again, PRS individuates mental qualities in a relational or holistic way—in terms of their relative locations within spaces that match corresponding spaces of perceptible properties. But many claim that in introspection mental qualities appear to be intrinsic properties. Recall, many traditional qualia theorists claim that mental qualities are, perhaps by definition, intrinsic properties. If mental qualities are intrinsic, holistic representationalism is a nonstarter.

The source of this introspective report, I think, is the fact that when we introspect our perceptions and sensations, we are typically aware of them only in respect of their particular mental qualities, not in respect of the connections between their particular qualities and other

[115] In developing responses to these objections, I have benefitted greatly from discussions with Prinz and Rosenthal.

qualities. Introspection is in that way atomistic. But though it is arguable that introspection does not present mental qualities as relational, it is also arguable that introspection does not present them as intrinsic. Rather, introspection is simply silent on the issue of whether mental qualities are intrinsic or relational.

Moreover, even if mental qualities really do appear intrinsic introspectively, introspection reveals only how mental qualities seem—and introspection need not always be accurate. The only reason to assume introspection is always accurate is the Cartesian assumption that the mind is always as it appears, an assumption which is dubious. As many psychologists and philosophers have argued, introspection is surprisingly limited and often mischaracterizes what is going on in one's mind (see, e.g., Hurlburt and Schwitzgebel 2007, Schwitzgebel 2008). Consider, for example, Richard Nisbett and Timothy Wilson's (1977) classic work, which nicely illustrates the unreliability of introspection. Through a series of ingenious experiments, Nisbett and Wilson demonstrated that people often confabulate, for example, the justifications for their preferences. Since introspective access to one's preferences is often fallible, it is likewise open that introspection may wrongly describe mental qualities as intrinsic when in fact such qualities are relational.

§5.6.2 There Are No Perceptible Properties

But there appears to be an even deeper worry for holistic representationalism. As several theorists have observed, there appears to be no interesting way in which the structure of the perceived similarities and differences between perceptible properties such as the perceptible colors matches the physical similarities that characterize those properties. In the case of perceptible color, this is often illustrated by the phenomena of so-called metamers—members of

a metameric pair of spectral reflectance properties are perceptually indistinguishable under particular illumination conditions, and yet are different physical properties. C. L. Hardin writes,

> There is, in fact, no single illuminant that will suffice to distinguish the members of every metameric pairs…. For now, we need only observe that there is little physical justification for putting [one member of the pair]… into the same color class with [the other member of the pair] (1993, p. 64, emphasis his).

On this basis of this and other evidence, Hardin observes that the "structure of the colors as we experience them corresponds to no known physical structure lying outside of nervous systems that is causally involved in the perception of color" (1997, p. 300, n. 2). On the basis of this evidence, many, including Hardin himself, have concluded that color realism—the view that colors are mind-independent and capable of being identified with physical attributes such as surface-reflectance properties—is false. Hardin writes that these facts make "it very difficult to subscribe to a color realism that is supposed to be about red, yellow, green, blue, black, and white—that is, the colors with which we are perceptually acquainted" (1997, p. 300, n. 2; for discussion of this critique, see, e.g., Byrne and Hilbert 2003, p. 13).

This evidence may seem to pose a serious problem for PRS. Of course, if there were in fact no perceptible colors, then the space of mental colors could not match the space of perceptible colors. But even if perceptible colors did exist, if there is no match between the structure of the perceived similarities and differences between the perceptible colors and the physical similarities that characterize those physical colors, then it may also seem that the space of mental colors cannot match the space of perceptible colors.

But there are several things to say about this kind of evidence. First, PRS is not a theory of the perceptible colors; it is a theory of mental qualities that explains them in terms of their relationship to perceptible colors. It is therefore beyond the purview of this dissertation to launch a full defense of color realism. But I will briefly sketch one credible realist option.

Recall, the human color-vision system operates according to an opponent-processing model, wherein three types of color-sensitive cones in the human eye are activated at different levels to produce visual sensations. According to this model, the reason why metamers are perceptually indistinguishable is that they activate equivalent ratios of these color sensitive cells. It is therefore open to identify the perceptible colors with equivalence classes of spectral reflectances that reflect distributions of wavelengths resulting in particular ratios of activation among the three color-sensitive cones in the eye under standard lighting conditions (see, e.g., Rosenthal 2005, p. 141; for a related view, see Byrne and Hilbert 2003, p. 12).

This proposal might seem to suggest that perceptible colors are not mind-independent properties. But while we fix the referents of our color terms by reference to the human visual system, this does not entail that perceptible colors are mind-dependent. As David Hilbert and Mark Kalderon put it:

> [T]he relations of similarity in hue, saturation, and brightness supervene on color properties that exist independently of viewing subjects. In speaking of the visual system of selecting these relations, we are not claiming that visual experience brings them into being. We are merely claiming that a pre-existing relation only counts as similarity in hue, saturation, or brightness in virtue of an antecedent classificatory function of the visual system. Facts about visual experience fix which of the similarities among certain objective properties count as similarities with respect to color. But that is consistent with higher-order color similarities, so understood, supervening on objective color properties. The properties represented by color experience are ontologically independent of those experiences, as are the relations of similarity and difference among them (2000, p. 198; cf. Rosenthal 2005, p. 197).

In other words, the perceptible colors are mind-independent properties of objects, but those mind-independent properties are grouped together into classes relative to the capacities of human perceivers. Those classes would exist even if there were no human beings, but we only delineate those classes as classes because of human capacities. There is no contradiction here. Even

though PRS is not a theory of the perceptible colors, there is no good reason to deny that the perceptible colors exist.

It is also important to note that holistic representationalism does not explain mental qualities in terms of their relationships to perceptible properties characterized in physical terms. Rather, holistic representationalism explains mental qualities in terms of their relationship to perceptible properties as those perceptible properties are discriminated. Like any property, perceptible properties can be grouped and characterized in many ways. We can, of course, characterize perceptible properties in terms of their physical characteristics. In the case of perceptible colors, we can characterize perceptible red, for example, as being any surface that reflects light with wavelengths in roughly the 630-740 nm range. Describing perceptible properties in this physical way, we can distinguish many different shades of perceptible red. We might call the shade that reflects wavelengths of 630 nm perceptible red-630, and distinguish that from perceptible red-700, which reflects wavelengths of 700 nm.

We can also characterize perceptible properties in terms of a creature's ability to discriminate them. On that way of grouping perceptible colors, perceptible red is the property that has a certain location in a quality space determined by a creature's discriminations—and for some creatures perceptible red-630 and perceptible red-700 will both be located at that position. Characterized in terms of those creatures' discriminatory abilities, there is no way to distinguish those two shades. Characterized physically, however, we are able to distinguish them. The space that orders perceptible colors according to their physical characteristics need not be, and is in fact not, identical to the space of perceptible colors ordered according to how they are discriminated. Two physical dimensions that are not contiguous in the space of perceptible

colors ordered physically may be located at the same position in the space of perceptible colors ordered by discriminations (i.e. a metameric pair).

PRS explains qualitative contents in terms of their locations in spaces that match the spaces of perceptible properties ordered by discriminations. Relative to the space of perceptible properties organized in that manner, the quality space of mental qualities is isomorphic to the space—that is, there is a one-to-one mapping of mental qualities to perceptible properties—considered as those perceptible properties are discriminated. By contrast, the quality space of mental qualities is homomorphic to that space—that is, there is a one-to-many mapping of mental qualities to perceptible properties—considered in respect of their physical characteristics.

§5.6.3 Undetectable Quality Inversion

Setting that objection aside, another objection to holistic representationalism might immediately spring to mind: Does not the conceivability of quality inversion entail that the spaces of mental qualities cannot match the spaces of perceptible properties? As noted in chapters 1 and 2, such a possibility is often cited as a reason to think that standard representationalism is true. The kind of holistic representationalism developed here is a variety of standard representationalism, except for the fact that its semantics of qualitative content is holistic.

But there are a couple of things to note here. First, it is often observed that quality inversion would be possible only if the spaces of mental qualities were symmetrical (e.g., Clark 1985, p. 431; Byrne 2010, §2.3.1). Since mental qualities are distinguished in terms of their unique locations in quality spaces, it would be impossible to distinguish one mental quality from another solely if the space of mental qualities were symmetrical and one quality resided on one

side of the axis of symmetry and the other quality was located at the corresponding position on the other side of the axis. Shoemaker observes,

> If behaviorally undetectable spectrum inversion is a possibility, then different color qualia may be in a certain sense functionally indistinguishable. They will occupy different locations in a "space" of qualia, but because of the symmetrical nature of that space it will be impossible to give a functional description that applies to the one but not to the other—rather in the way that if one is describing in spatial terms a spatially symmetrical array of objects, then any description one gives of an object will apply as well to its symmetrical counterpart. If this is so, then color qualia will be "ineffable" in a way that goes beyond the irreducibility to particular physical properties that goes with multiple realizability (1994a, pp. 306-307).

Since PRS posits that spaces of mental qualities are extrapolated from the spaces of perceptible properties, the spaces of mental qualities would be symmetrical only if the corresponding spaces of perceptible properties were symmetrical. It is the case, however, that all of the quality spaces of perceptible properties that have been mapped are asymmetrical. For example, discussing asymmetries in the space of perceptible colors for human visual systems, Byrne observes,

> [T]here are more perceptually distinguishable shades between red and blue than there are between green and yellow, which would make red-green inversion behaviorally detectable. And there are yet further asymmetries. Dark yellow is brown (qualitatively different from yellow), whereas dark blue is blue…. Similarly, desaturated bluish-red is pink (qualitatively different from saturated bluish-red), whereas desaturated greenish-yellow is similar to saturated greenish-yellow (2010, §2.3.1).

One might think that there are, nevertheless, some quality spaces that are symmetrical—perhaps one-dimensional spaces such as the space of achromatic colors (viz. white to black) that differ only in terms of brightness. But there is evidence that even such one-dimensional spaces exhibit asymmetry. For example, people typically perceive the brightest element of an achromatic series to be white, regardless of its physical magnitude (e.g., Gilchrist 2009). In this way, white anchors the one-dimensional space, rendering it asymmetrical (see, e.g., Rosenthal 2005, p. 225).

Even so, some argue that the fact that actual perceptible-color spaces are asymmetrical is of no consequence, because it is conceivable that they might not be. Shoemaker writes,

> Even if our color experience is not invertible, it seems obviously possible that there should be creatures, otherwise very much like ourselves, whose color experience does have a structure that allows for such a mapping—creatures whose color experience is invertible. And the mere possibility of such creatures is sufficient to raise the philosophical problems the possibility of spectrum inversion has been seen as posing (1982, p. 367).

But, as many representationalists have argued, inversion is not possible and arguably inconceivable (e.g., Harman 1996, p. 12; Speaks 2011). As noted previously, mental qualities arguably appear intrinsic in introspection, which explains why it may seem to many that such inversion is possible. But the actual possibility of quality inversion depends on the groundless assumption that mental qualities are only known by way of consciousness. I've argued this assumption is false and that it is more coherent to conceive of mental qualities in terms of their representational roles in perception. The alternative theory, holistic representationalism, tells us what mental qualities are. Without such a theory, many aspects of mental qualities may seem possible that are in fact impossible. Indeed, in the absence of a suitable theory about some phenomenon, many things may seem possible about it. For instance, in the absence of a suitable theory of light, it might have seemed possible that light could have traveled at 500,000 miles per second, instead 186,282 miles per second. But with a theory of light in hand, it is impossible that light could have traveled at the faster speed. Likewise, once we have an informative account of mental qualities, it is clear that quality inversion is impossible.[116]

[116] Shoemaker's view of mental qualities already diverges from the present account insofar as he claims that perception represents appearance properties, not ordinary perceptible properties (see chapter 3). But by holding that quality inversion is conceivable, Shoemaker's account of mental qualities also rejects what I see as another fundamental plank of PRS and holistic representationalism—namely, that mental qualities are individuated by their unique locations within quality spaces. Since Shoemaker assumes that the quality spaces of mental qualities are (or at least could be) symmetrical, he holds that such quality spaces do not individuate mental qualities. It is therefore unclear what Shoemaker thinks individuates them. In any case, as I've argued, Shoemaker provides no good reason to think that the spaces could be symmetrical and hence why we cannot individuate mental qualities by their unique locations in those spaces.

§5.6.4 Nonconceptual Content and Nonrepresentational Qualia

According to holistic representationalism, the nonconceptual contents that PRS posits should be identified with mental qualities. But some proponents of nonconceptual content are explicit that nonconceptual contents are distinct from mental qualities. So there is a question whether holistic representationalism's claim that mental qualities are a species of nonconceptual content is licit.

Christopher Peacocke (1992), for example, clearly distinguishes nonrepresentational qualia, which he calls sensational properties, from two sorts of nonconceptual content that he also posits. According to Peacocke, sensational properties are the properties in virtue of which sensations differ qualitatively. A sensation of red, for example, differs from a sensation of green because the former exhibits a red sensational property, whereas the latter exhibits a green sensational property. Importantly, Peacocke argues that sensational properties correspond to, but do not represent, perceptible properties.[117] Peacocke posits sensational properties to account for a variety of phenomena, including quality inversion and perceptual constancies. According to Peacocke's (1983, p. 12) oft-cited example involving size constancy, two trees can be visually represented as being the same height, even though the closer object occupies a larger portion of one's visual field. Since the trees are not represented both as being the same height and as being different heights, Peacocke argues that difference between the perceptions a difference in their nonrepresentational sensational properties, not in their representational contents.

In addition to sensational properties, Peacocke posits a form of nonconceptual content, which he calls positioned scenario content, which he claims "should be individuated by specifying which ways of filling out the space around the perceiver are consistent with the

[117] Peacocke distinguishes sensational properties from their corresponding perceptible properties by using primed predicates. For example, red' is the sensational property of sensations of perceptible red.

representational content's being correct" (1992, p. 61). This kind of content represents what Peacocke calls a scene, which is a way of locating perceptible properties in relation to an origin and family of axes centered at the perceiver's location (1992, p. 64). Peacocke writes, "For each point… identified by its distance and direction from the origin, we need to specify whether there is a surface there and, if so, what texture, hue, saturation, and brightness it has at that point" (1992, p. 63). These properties and relations in the scene, Peacocke claims, are represented nonconceptually by positioned scenario content.

Peacocke also posits another kind of nonconceptual content, which he calls protopropositional content. Peacocke posits this kind of content to explain how we perceive certain objects, properties, and relations in ways that may also outstrip our conceptual capacities, but that do not involve filling out a scene (and therefore do not involve positioned scenario content). For example, a perception of a square differs representationally from a perception of a diamond, but the perceptions arguably do not differ in terms of how one fills out the space around the perceiver. Thus Peacocke argues the difference lies at the level of their nonconceptual protopropositional content (see, e.g., 1992, pp. 74-77).

Peacocke offers several arguments for thinking that these two varieties of content are nonconceptual, including arguments discussed in chapter 2. Though these arguments carry some weight, they do not establish that perception exhibits nonconceptual content and nonrepresentational qualia as against the view that perception exhibits representational mental qualities only. So why prefer the present view to Peacocke's account? One problem for Peacocke's account is that it posits too much. Representational mental qualities alone can explain what Peacocke posits both kinds of nonconceptual content and nonrepresentational sensational properties to explain (see, e.g., Meehan 2002). As I've argued in this chapter, we

need not posit nonrepresentational sensational properties to explain perceptual constancies. Such effects can be explained by representational mental qualities. One's experience veridically represents the objects as being the same height in virtue of exhibiting conceptual representational content but veridically represents them via size mental qualities as having size properties in virtue of which they subtend different retinal angles. Likewise, the perceptual representations that Peacocke posits his two kinds of nonconceptual contents to explain can be explained by representational mental qualities too. Positioned scenario contents concern ordinary perceptible properties and the difference between the perceptions of, for example, a square and a diamond can be explained by differences in their spatial mental qualities.

But considerations of simplicity are not the only motivations for preferring the account developed here to Peacocke's. Perhaps more importantly, because Peacocke posits sensational properties largely to explain undetectable quality inversion, Peacocke cannot offer any informative account of them. As I argued in chapter 3, three of the most developed accounts of mental qualities that attempt to accommodate undetectable quality inversion cannot provide informative accounts of mental qualities. Peacocke's account cannot do better.

§5.6.5 The Disjunction Problem

There are also a variety of objections to a host of theories of conceptual representation, and one might think that these apply to the account of qualitative representation defended here. For example, Jerry Fodor has maintained that any theory of representation must solve what he terms the disjunction problem (1987, p. 102). Fodor raises the disjunction problem as an objection to simple causal-covariational accounts of intentionality of the sort developed by Dretske (1981). According to such accounts, a concept represents something just in case, roughly, that thing

reliably causes the concept to token. But Fodor observes that there are examples of things' reliably causing certain concepts to token, even though those concepts do not refer to those things. Because cleverly disguised cats reliably cause the concept DOG to token, a causal-covariational theory seems to have the upshot that the concept DOG represents the disjunction of dogs and cleverly disguised cats. But that consequence is wrong: DOG refers to dogs only.

Whether or not the disjunction problem is a problem for some accounts of conceptual representation, one might think that an analogue of it plagues the present account of qualitative representation. For example, one might wonder why mental red represents perceptible red and not the disjunction of perceptible red and perceptible blue. No version of the disjunction problem is, however, a problem for PRS. According to PRS, there is a ready explanation of why a particular mental quality represents the particular perceptible property that it does (as opposed to any disjunctive properties). Unlike simplistic causal-covariational accounts of representation, PRS holds that mental qualities' representational characters are determined by their locations within quality spaces that match corresponding spaces of perceptible properties. Because the spaces of mental qualities are extrapolated from those spaces of perceptible properties, the one-to-one correspondence of mental qualities to perceptible properties is guaranteed.

§5.6.6 Ganzfelds

So far, the objections above target most versions of strong representationalism. In his recent book, however, Prinz (2012) offers some objections to this particularly holistic version of representationalism.[118] I will consider these in turn.

[118] Prinz's target is the view that he calls 'experiential holism', which I take to capture the spirit of holistic representationalism. Prinz attributes experiential holism to, among others, Clark (1993) and Rosenthal (2005).

First, Prinz offers a commonsense reason to doubt holistic representationalism. Prinz claims that qualitative content cannot be individuated in a holistic way because an "experience has its phenomenal character even when no other experiences of the same kind are occurring" (2012, p. 132). Prinz imagines here that one may have an experience of only a single color, for example a so-called ganzfeld consisting solely of red, and that this possibility suggests that the state's qualitative content cannot be determined by its relationship to other qualities.

To see why this objection does not go through, consider an analogous objection to a holistic account of conceptual content such as CRS. It is, of course, true that one can think a single thought at a given time. One can think, for example, that New York is the Empire State, even if one is not also currently thinking other thoughts. But it would be illegitimate to argue that CRS is false because one is able to think only a single thought at a particular time. Though CRS characterizes the contents of thoughts in terms of their inferential relations to one another, it does not hold that one must currently engage in inferences or have multiple thoughts at a time to exhibit a thought. Rather, CRS characterizes the contents of thoughts in terms of one's dispositions to draw inferences.

Likewise, PRS characterizes qualitative contents in terms of their similarity-and-difference relations to one another. But these are not relations to other qualitative states that one need exhibit occurently, but rather perceptions that one can exhibit. Spaces of mental qualities are not themselves represented in the mind. Spaces of mental qualities are representations of an individual's perceptual discriminatory abilities, whether or not the creature is engaging in perceptual discrimination at any given time. PRS thus holds that a qualitative content exhibits a location within a space whether or not that content is currently tokened alongside other contents

that make up that space. One need not exhibit a mental quality in order for it to play a role in determining the content of the qualitative states that one does exhibit.

For Prinz's objection to undermine PRS, it would have to be the case not only that a creature perceive only a single color at a time, but be capable of experiencing only a single color. But it is not obvious that such a creature is possible. Indeed, there is experimental evidence that a creature exposed to a ganzfeld only will in time lose the capacity to see anything (e.g., Metzger 1930; for discussion, see, e.g., Gibson 1979, pp. 150-151). While this does not entail that the contents of perceptions are solely determined by their relations to other currently exhibited perceptions, it does suggest, as PRS holds, that qualitative content is relational. Thus it is perfectly consistent with holistic representationalism that one may perceive only a single color at a time, though one's capacity to perceive that color depends on one's ability to perceptually distinguish among a repertoire of colors.

§5.6.7 No Qualities in Common

Prinz also offers a theoretical reason to doubt holistic representationalism, invoking the well-known evidence that there is great diversity among individuals' perceptual capacities. Since PRS holds that a qualitative state's content—and thus its qualitative character—is determined by its unique location within a space of qualitative contents extrapolated from the quality spaces of perceptible properties that a creature can perceptually distinguish, Prinz argues that such holism entails that no two individuals will be likely to exhibit the same qualitative character because their respective spaces of perceptions will differ from one another. As Prinz puts it, a consequence of holistic representationalism is that "two individuals with a different set of color representations will not have any color experiences in common, because no pair of color

representations in these two systems will have the same role" (2012, p. 132). And Prinz argues that this consequence is not credible insofar as it is intuitively obvious that two individuals can exhibit perceptions with the same qualitative character.

Prinz's argument here echoes a common challenge to holist theories of conceptual content such as CRS. Since CRS individuates thoughts' contents in terms of their relative locations within networks of inferential relations, critics of the theory have observed that no two individuals will be likely to have thoughts with the same contents since any two individuals are likely to draw different sets of inferences (e.g., Fodor and Lepore 1992, pp. 17-22). And this would seem to pose a grave problem for CRS: If two people cannot exhibit thoughts with the same content, then it is unclear, for example, how meaningful communication can take place between those two individuals.

Whatever the efficacy of this critique of CRS, the analogous objection to PRS does not pose a problem for the view. Even if it is true that any two individuals will exhibit slightly different perceptual discriminatory capacities—and thus exhibit slightly different spaces of mental qualities—it does not follow that such individuals cannot exhibit roughly the same qualities. To the extent that the two individuals' spaces of perceptions are roughly the same, differences between their respective perceptual discriminations will not be meaningful for any functional purpose. Thus, even if it is true that the two people do not exhibit exactly the same mental qualities, there would be no reason to expect that ordinary people would attribute to them differing mental qualities. Indeed, common sense only attributes distinct mental qualities to people when they make radically different perceptual discriminations from the norm, such as in cases of colorblindness.

§5.6.8 Abnormal Vision

Prinz offers in addition experimental evidence that seems to question holistic representationalism. Prinz cites evidence from Lothar Spillmann and colleagues (2000) regarding a patient who had suffered a brain injury resulting in a disorder known as color agnosia, a form of colorblindness. As Prinz reports, the color sensitivity of Spillmann and colleagues' patient returned gradually, color by color, rather than all at once. The patient first was able to recognize perceptible green, then perceptible red, and then perceptible blue. This case seems to suggest that perceptions' contents cannot be identified holistically.

But there are two ways to respond to this evidence. First, recall that PRS holds that a qualitative state's content is determined by its location within a space of mental qualities, whether or not that state is currently tokened alongside the other states that make up that space. So if Spillmann's patient's perceptual discriminatory capacities remained intact throughout his recovery, then his perceptions' contents can be individuated in a holistic way even if he lacks the capacity to token some perceptions at any given time.

Alternatively, one can take issue with how the case is described. Spillmann's patient is described as being capable of discriminating first only perceptible green things, then perceptible red things, and then perceptible blue things. But just because the patient was able initially to discriminate perceptible green things from their surrounds does not show that the patient was qualitatively representing those green things as perceptible green—that is, as exhibiting perceptions with mental green. It could be instead that he initially misrepresented those perceptible green things as some other perceptible color, which enabled him to distinguish them. Thus it is arguable that the character of his perceptions changed as he gained the capacity to discriminate more colors, as PRS would predict.

§5.6.9 Shifted Quality Spaces

Prinz (personal communication) suggests yet another objection to holistic representationalism. Prinz asks us to imagine a creature with a visual system different from ours, one that exhibits a "shifted spectrum" such that, for every location within that imagined creature's quality space of perceptible colors, the wavelengths located at each location are shifted from the space of perceptible colors generated by normal human adults. Suppose, for example, that actual human adults discriminate surfaces that reflect X wavelengths of light as being located within the quality space of perceptible colors at location L. This imagined creatures, by contrast, discriminates surfaces that reflect X+1 wavelength of light as being at location L in its space of perceptible colors. And the same goes for any property located at any location within the quality space.

Of course, such a creature would be behaviorally distinguishable from normal humans. But since PRS individuates spaces of mental qualities relative to spaces of perceptible properties understood as those properties are discriminated, not as those properties are construed physically, such an imagined creature's quality space of perceptible colors would exhibit the same geometrical properties as the normal human space of perceptible colors. In other words, the creature would regard a different set of stimuli understood physically as standing in the same similarity-and-difference relations as stand the set of perceptible colors (as distinguished by normal human adults).

The problem, Prinz argues, is that PRS holds that mental qualities are individuated by their locations within spaces that match the spaces of perceptible properties construed in terms of how those properties are discriminated. Thus PRS must hold that such an imagined creature has the same space of mental colors as normal humans. This is not credible, however, because each

property construed physically at any location in the imagined creature's space of perceptible properties is distinct from the property construed physically at the corresponding location in the normal human's space of perceptible properties. So, Prinz asks, does the mental quality sitting at the location in the quality space of mental colors corresponding to location L in the space of perceptible colors represent property X or property X+1? It would seem PRS must say it corresponds to both.

One might think that, since perceptible colors are already characterized in terms of equivalence classes, one could simply hold that X and X+1 are members of the same class. But Prinz argues that this reply opens PRS up to a version of the sorites paradox. One can also imagine creatures that discriminate spaces of perceptible properties wherein property X+2 is located at L, X+3 is located at L, and so on. If such creatures are possible, then any perceptible property (construed physically) potentially could be located at any location within the quality space of perceptible colors. But if there is no non-arbitrary way to distinguish the equivalence classes that characterize each perceptible property, then there would be no way to determine to which property each mental quality corresponds—and it appears that PRS is sunk.

But, again, PRS is not a theory of perceptible properties such as colors. Thus it is not incumbent upon the quality-space theorist to explain what distinguishes the equivalence classes that characterize the perceptible colors. Distinguishing such classes is a job for the science and philosophy of color. Whatever the way to distinguish those classes, creatures that discriminate "shifted" spaces may not count as discriminating among the perceptible colors. Though such imaginary creatures discriminate spaces of properties that have the same geometrical features as the spaces of perceptible colors, such spaces are not the space of the colors. Instead, we might say that such creatures exhibit a distinct kind of sensory modality sensitive to another kind of

perceptible property. That is, it is open to hold that such creatures perceive perceptible properties akin to perceptible colors—which we might call colors*, colors**, and so on—via distinct perceptual modalities—which we might call (because of their similarities to vision) vision*, vision**, and so on. Such a reply may sound extreme, but recall that such creatures would be very different behaviorally than normal humans.

§5.6.10 Reorganized Quality Spaces

A similar objection involves imagining creatures that exhibit not shifted, but reorganized spaces of perceptible properties. Again, imagine there are creatures with different visual systems whose distinct discriminatory capacities generate spaces of properties that have the same geometrical properties as the space of perceptible colors discriminated by normal humans. But, because of their different perceptual systems, such creatures' spaces exhibit different organizations of the same set of properties construed physically. So, while the creatures can discriminate the same set of properties, they judge them to stand in very different similarity-and-difference relations.

Such creatures would be, of course, quite behaviorally distinguishable from normal humans and from one another. But an analogue of the problem raised by Prinz above arises here. Since any property (construed physically) can be located at any location within the space, then again, there would seem to be no way to determine which property each mental quality corresponds to in the matching space of mental qualities—and it appears that PRS cannot work.

But the same rebuttal to the objection concerning shifted quality spaces applies here as well. Given how behaviorally distinct such creatures would be from ordinary humans, it is hard to see why we should regard them as sensing the perceptible properties via the same perceptual modalities as we do. Because such creatures would be distinguishing among the same set of

perceptible properties construed physically (e.g., the perceptible colors), we might hold that they perceive the perceptible colors. But since such creatures make such wildly distinct discriminations from normal humans, there is reason to think they perceive those properties via distinct sensory modalities. Such a situation would be akin to the way in which distinct sensory modalities in normal humans detect the same perceptible properties (for example, as vision and touch can both detect perceptible shapes). That is, such creatures would be accessing the perceptible colors via vision*, vision**, and so on, instead of vision, which results in their reorganizing the perceptible colors construed physically.

§5.6.11 Categorical Perception

Before closing this section, I'll discuss one last piece of experimental evidence that may seem to undermine holistic representationalism. There is much evidence of what is known as categorical perception—namely, that some perceptible properties are perceived not along a smooth continuum, but instead are perceived as belonging to discrete categories (see, e.g., Harnad 1987). For instance, most normal human perceivers draw a relatively sharp distinction between shades of perceptible green and shades of perceptible blue. And there is growing evidence that such categorical distinctions between perceptible colors are susceptible to cultural influences (for a review, see, e.g., Goldstone and Hendrickson 2010). It is well-known that different cultures have different color categories, as is evidenced by their different color vocabularies. For example, Russian speakers make an obligatory distinction between lighter blues, which they call 'goluboy', and darker blues, which they call 'siniy'. English, by contrast, does not make this distinction. Interestingly, Jonathan Winawer and colleagues (2007) have shown experimentally that Russian speakers are faster at drawing distinctions across the categorical divide of light blue

shades and dark blue shades than their English-speaking counterparts, suggesting that learned color categories can affect how one categorically perceives colors.

This evidence might seem to suggest a problem for holistic representationalism. If Russian speakers perceive a sharp distinction between light and dark blues, then perhaps such perceptible colors cannot fall into a continuous color quality space—and so cannot in turn determine those perceivers' spaces of mental qualities. But quality spaces simply are reflections of creatures' abilities to make perceptual distinctions among some family of perceptible properties. If a creature makes such perceptual discriminations, its behavior can be modeled by a quality space. And if, for example, those perceptual discriminations involve sharp distinctions among certain qualities, the space can capture this fact by increasing or decreasing the distance between qualities. So, for example, if Russian speakers draw a distinction between light blue and dark blue, this can be captured by decreasing the distance between light blue and dark blue.

But one might also think that the fact that learned categories can affect how one categorically perceives perceptible properties undermines PRS. But though PRS is arguably committed to the view that some quality space(s) must be innate, it is neutral whether such spaces can be adjusted by learning or other developmental processes. PRS holds that the quality spaces which determine one's qualitative contents supervene on one's occurrent perceptual discriminatory capacities. Whether those perceptual discriminatory capacities can be adjusted due to various influences is an open empirical question.

§5.7 Applying the Account to Other Debates in the Philosophy of Mind

I'll now discuss how this view regarding the representational nature of mental qualities can be brought to bear on several prominent debates within contemporary philosophy of mind.

§5.7.1 The Contents of Perception

Recently, there has been considerable debate regarding the so-called contents of perception, a debate concerning which properties are represented by perceptions and which properties are represented in downstream thoughts (e.g., Siegel 2006, Siegel 2010a). As noted in chapter 1, most theorists today assume that perception is representational. And those who hold that perception represents things do generally agree that perception represents perceptible properties such as colors and smells. But there is question whether or not a perception also represents higher-level properties such as the property of being a chair. Attributions of perceptions often involve references to such higher-level properties, as in Jack perceives that there is a chair, but some deny that the perception thereby represents those higher-level properties. For example, representationalists such as Dretske (1995) and Tye (1995) maintain that perception represents low-level perceptible properties only. Susanna Siegel, by contrast, argues that perception does represent higher-level properties.

According to the view of perception developed here, perception involves two factors: Mental qualities that represent perceptible properties only and perceptual conceptual content. As Siegel observes, the debate over the contents of perception becomes considerably more tractable if one adopts this sort of two-factor view. Calling the view that perception has content at all the Content View and the view that perception represents higher-level properties (which she calls K-properties) the Rich Content View, Siegel writes,

> Neither the Content View nor the Rich Content View is very interesting when applied to visual experiences of seeing-as that are structured by two factors as described above. The Content View would be true of those states, thanks to their representational factor, and the Rich Content View would be true if that factor ever represents K-properties. It is no surprise that judgments have contents, or that they can represent K-properties. For this reason, the Content View and the Rich Content View are theses about visual experiences in the narrow sense: visual phenomenal states. The Content View is the thesis that all

> visual phenomenal states have contents. The Rich Content View is the thesis that some visual phenomenal states represent K-properties (2010a, p. 22).

Siegel regards the debate concerning the contents of perception as genuinely puzzling because she approaches the phenomenon of perception from a one-factor position. I have argued here that the one-factor position is false. But Siegel, like most traditional qualia theorists, assumes that if a two-factor view involves mental qualities, then those qualities are not representational.

It is thus clear how PRS can help us make headway on the issue of the contents of perception. According to PRS, the properties represented in perception relevant to determining qualitative character are the perceptible properties that fall into well-defined families discriminated via a sensory modality. It is, however, an open empirical question which properties those are. If perception also exhibits conceptual content, then there is an open question whether those contents concern the perceptible properties also represented by mental qualities only, higher-level properties only, or a combination of both. I will not attempt to settle this issue now. But it is not clear what substantive issues turn on settling it. Additionally, since ordinary attributions of perceptions make reference to higher-level properties, there seems to be no good reason to deny that perceptual conceptual content concerns those properties.

§5.7.2 Enactive Perception

Another prominent debate in philosophy of mind concerns the view known as the enactive approach to perception (e.g., Hurley 1998, O'Regan & Noë 2001). Alva Noë (2004) presents the most well-developed version of this account, so I will focus on it. According to Noë, perception constitutively involves action. That is, Noë claims that in order to perceive something, one must bring to bear what Noë calls sensorimotor knowledge, implicit knowledge of how a thing will look depending on how one interacts with it (e.g., 2004, pp. 1-2).

Noë offers a range of evidence in support of this enactive approach to perception, but a major motivation for it is the commonsense observation that we typically see more of an object than is visually present to one at any given time. For example, one sees a coin as round, even though at a given time one is presented with only one side of the coin from one visual angle—say, when one has tilted it in such a way that it appears elliptical. Noë claims that to see a tilted coin as round (even when it appears elliptical) involves the exercise of a certain kind of implicit knowledge concerning how the coin's appearance would change if one were to turn it, move around it, and interact with it generally. On Noë's view, vision appears to provide us with a great deal of detail about our environments, when in fact we are only presented with a narrow range of features.

Noë takes the perspective-dependent ways that objects can appear to be relational properties of the objects, which he calls perspectival properties (2004, pp. 82-84). Like Shoemaker's appearance properties discussed in chapter 3, perspectival properties are related to but not identical with perceptible properties. Unlike Shoemaker's appearance properties, however, Noë claims that perspectival properties are mind-independent relational properties that hold between objects and vantage points. On Noë's view, one sees a tilted coin as exhibiting the perspective-invariant property of perceptible roundness because one knows that it would exhibit the perspective-dependent property of perspectival ellipticality if one were to set it on an angle.

At first blush, one might think that Noë's enactive approach might fit nicely with the account of qualitative representation developed here. However, Noë acknowledges that his account owes much to the work of J. J. Gibson (1979), who maintained that seeing the properties of an object does not involve visual representations of those properties. Rather, Gibson argues, as Noë glosses it, that "animals… are directly sensitive to the features of the world that afford the

animal opportunity for action (what Gibson 1979, chap. 8, called "affordances")" (2004, p. 21). I will not explore Gibson's theory in more detail. I mention it because, even though Noë does not go so far as to deny that there are any representations in perception,[119] his theory does seem to involve the view that perception does not operate in terms of qualitative representations of perceptible properties, as I have been arguing.

Importantly, Noë argues that we cannot distinguish the aspects of perception that concern the way an object currently appears and how it would appear were one to interact with it, on the one hand, from the way that the object is independent of how it appears at a given time. Because one's perception of the tilted coin as round constitutively involves an awareness of a complex array of facts concerning how the coin would look under various conditions, Noë claims that all of those features are somehow both present and absent in one's perception. That is, the content of one's perception of the coin is, as Noë puts it, "virtual" and that perceptual content is "virtual all the way down... [which] shows that we cannot factor experience into an occurrent a merely virtual or potential part" (2004, pp. 134-135, emphasis his). Noë's claims are a bit obscure here, but he goes on to clarify that "[q]ualities are available in experience as possibilities, as potentialities, but not as givens" (2004, p. 135). On the view defended here, by contrast, one does represent perceptible properties, not as givens (whatever that amounts to), but also not only as possibilities. So it would seem that Noë's proposal is actually quite radical and does challenge the view developed in this chapter. Indeed, Noë outright admits that "the qualia theory is not compatible with an enactive account of color experience" (2004, p. 124).

Noë's account does not, however, successfully undermine the account of qualitative representation developed here. As noted before, Noë's enactive approach to perception depends

[119] Noë claims that his enactive approach entails only that "the role of representations in perceptual theory needs to be reconsidered" (2004, p. 22).

crucially upon the perception of perspectival properties. Seeing a coin as exhibiting the perspective-invariant property of perceptible roundness involves understanding what perspectival properties one would perceive the coin to exhibit were one to interact with the coin in various other ways, such as perceiving the coin to exhibit perspectival ellipticality if one were to set the coin at a tilted angle. In this chapter, I have defended a view of perception that adequately accounts for the nature of perception without positing perspectival properties in addition to perceptible properties. On the view developed here, we can explain why the coin both looks elliptical and round by holding that we may represent those perceptible properties in two different ways—that is, we conceptually represent the coin as exhibiting perceptible roundness and qualitatively represent perceptible ellipticality.

But even if Noë were correct that objects exhibit perspectival properties, it is not clear how his enactive account can explain how we perceive those properties (e.g., Meehan ms., pp. 126-128). Noë claims that the perception of a perspectival property constitutively involves recruiting one's sensorimotor knowledge about it (e.g., 2004, p. 90), but it is not clear what that knowledge consists in or how one is supposed to apply it. Sensorimotor knowledge involves knowing what perspectival properties an object would exhibit if one were to interact with it, but perspectival properties themselves do not exhibit perspectival properties. Indeed, it is a mystery how one might interact with perspectival ellipticality such that one could have knowledge of how that appearance itself appears, were one to interact with it. Moreover, it would seem that one must first see a perspectival property in order to interact with an object that exhibits it.[120]

[120] In a review of Noë's book, Michael Pace (2005) makes a similar critique. Discussing the expectations involved in the application of sensorimotor knowledge during perception, Pace writes, "I cannot see how to make sense of an expectation [that one will see something given some action] without postulating an occurrent representation as to how things appear now and a

Noë thus cannot explain how his enactive approach gets off the ground if we do not represent some properties in a manner that is not enactive. I will not adjudicate whether Noë is correct that we perceive the perspective-invariant perceptible properties of an object such as an object's perceptible roundness in an enactive way, though several have criticized the theory on various other grounds (e.g., Prinz 2006). But for Noë's enactive account to work at all, the theory must explain how we are perceptually aware of either perceptible or perspectival properties. A straightforward account of these matters is the one developed here—that we qualitatively represent such properties.

§5.8 Conclusions

Representationalism about qualitative character is very attractive. But most versions of representationalism depend on simple atomistic theory of the contents of qualitative states—and I've argued that these versions of representationalism are problematic. There is overwhelming evidence that qualitative character is, at bottom, a holistic or comparative phenomenon. Fortunately, a holistic semantics for qualitative content is available—the view that I have called PRS. And I've argued that this resultant holistic version of representationalism is the best view of qualitative character currently available.

---

representation of the way movements would change the appearances, both of which seem to be internal matters."

References

Adams, F. and L. Dietrich (2004), "Swampman's Revenge: Squabbles Among the Representationalists," Philosophical Psychology 17, 3: 323-340.

Albers, J. (2006), Interaction of Color, revised and expanded edition, New Haven: Yale University Press.

Aristotle (1968), de Anima: Books II, III (with passages from Book I), tr. D. W. Hamlyn, Oxford: Clarendon Press.

Armstrong, D. M. (1968), A Materialist Theory of the Mind, London: Routledge and Kegan Paul.

Austin, J. L. (1962a), Sense and Sensibilia, ed. G. J. Warnock, Oxford: Oxford University Press.

--- (1962b), How to Do Things with Words, 2nd edition, eds. J. O. Urmson and M. Sbisà, Cambridge, MA: Harvard University Press.

Aydede, M. and G. Güzeldere (2005), "Cognitive Architecture, Concepts, and Introspection: An Information-theoretic Solution to the Problem of Phenomenal Consciousness," Noûs 35, 2: 197--255.

Ayer, A. J. (1956), The Problem of Knowledge, London: Macmillan.

Baars, B. J. (1988), A Cognitive Theory of Consciousness, Cambridge, MA: Cambridge University Press.

Balog, K. (2009), "Phenomenal Concepts," in Oxford Handbook to the Philosophy of Mind, eds. A. Beckermann, B. P. McLaughlin and S. Walter, Oxford: Oxford University Press, pp. 292-312.

--- (2012), "Acquaintance and the Mind-Body Problem," in The Mental, the Physical, eds. C. Hill and S. Gozzano, Cambridge, UK: Cambridge University Press, pp. 16-42.

Bengson, J., E. Grube, and D. Z. Korman (2011), "A New Framework for Conceptualism," Noûs 45, 1: 167-189.

Berger, J. (2012), "Do We Conceptualize Every Color We Consciously Discriminate?" Consciousness and Cognition 21, 2: 632-635.

Berkeley, G. (1710/1996), Principles of Human Knowledge and Three Dialogues, ed. H. Robinson, Oxford: Oxford University Press.

Bermúdez, J. L. (1998), The Paradox of Self-Consciousness, Cambridge, MA: MIT Press.

--- (2007), "What is at Stake in the Debate on Nonconceptual Content?" Philosophical Perspectives 21: 55-72.

Block, N. (1978), "Troubles with Functionalism," Minnesota Studies in the Philosophy of Science IX, 261-325; reprinted in Block (2007a), pp. 63-101.

--- (1986), "Advertisement for a Semantics for Psychology," Midwest Studies in Philosophy 10, 1: 615-678.

--- (1990), "Inverted Earth," Philosophical Perspectives 4, Action Theory and the Philosophy of Mind: 53-79; reprinted in Block (2007a), pp. 511-532.

--- (1995), "How Many Concepts of Consciousness?" Behavioral and Brain Sciences 18, 2: 272-287; reprinted in Block (2007a), pp. 215-247.

--- (1996), "Mental Paint and Mental Latex," Philosophical Issues 7: 19-49.

--- (1999), "Sexism, Racism, Ageism, and the Nature of Consciousness," Philosophical Topics 26, 1: 39-70; reprinted in Block (2007a), pp. 571-602.

--- (2003), "Mental Paint," in Reflections and Replies: Essays on the Philosophy of Tyler Burge, eds. M. Hahn and B. Ramberg, Cambridge, MA: MIT Press, pp. 165-200; reprinted in Block (2007a), pp. 533-570.

--- (2004), "Qualia," in The Oxford Companion to the Mind, 2nd edition, ed. R. L. Gregory, Oxford: Oxford University Press, pp. 785-789; reprinted in Block (2007a), pp. 502-510.

--- (2007a), Consciousness, Function, and Representation: Collected Papers, Volume 1, Cambridge, MA: MIT Press.

--- (2007b), "Wittgenstein and Qualia," Philosophical Perspectives 21, 1: 73-115.

--- (2007c), "Max Black's Objection to Mind-Body Identity," in Phenomenal Concepts and Phenomenal Knowledge: New Essays on Consciousness and Physicalism, eds. T. Alter and S. Walter, Oxford: Oxford University Press, pp. 249-306; reprinted in Block (2007a), pp. 535-498.

--- (2010), "Attention and Mental Paint," Philosophical Issues 20, 1: 23-63.

Block, N. and J. A. Fodor (1972), "What Psychological States are Not," Philosophical Review 81: 159-81.

Boghossian, P. A. and J. D. Velleman (1989), "Color as a Secondary Quality," Mind 98: 81-103.

Bourget, D. (2010), "Consciousness is Underived Intentionality," Noûs 44, 1: 32-58.

Bourget, D. and A. Mendelovici (forthcoming), "Tracking Representationalism: Lycan, Dretske, and Tye," in Philosophy of Mind: The Key Thinkers, ed. A. Bailey, Continuum.

Breitmeyer, B. G., T. Ro and N. S. Singhal (2004), "Unconscious Color Priming Occurs at Stimulus-Not Percept-Dependent Levels of Processing," Psychological Science 15, 3: 198-202.

Breitmeyer, B. G. and H. Öğmen (2000), "Recent Models and Findings in Visual Backward Masking: A Comparison, Review, and Update," Attention, Perception, and Psychophysics 62, 8: 1572-1595.

Brentano, F. (1874/1973), Psychology from an Empirical Standpoint, ed. L. McAlister, London: Routledge and Kegan Paul.

Brewer, B. (1999), Perception and Reason, Oxford: Oxford University Press.

--- (2006), "Perception and Content," The European Journal of Philosophy 14: 165--81.

Burge, T. (2003), "Qualia and Representational content: Reply to Block," in Reflections and Replies: Essays on the Philosophy of Tyler Burge, eds. M. Hahn and B. Ramberg, Cambridge, MA: MIT Press, pp. 405-415.

--- (2010), The Origins of Objectivity, Oxford: Oxford University Press.

Burnyeat, M. (1992), "Is An Aristotelian Philosophy of Mind Still Credible? (A Draft)," in Essays on Aristotle's De Anima, eds. M. Nussbaum and A. Rorty, Oxford: Oxford University Press, pp. 15-26.

Byrne, A. (2001), "Representationalism Defended," Philosophical Review 110, 2: 199-240.

--- (2005), "Perception and Conceptual Content," in Contemporary Debates in Epistemology, eds. M. Steup and E. Sosa, Malden, MA: Blackwell, pp. 231-250.

--- (2006), "Color and the Mind-Body Problem," Dialectica 60, 2: 223-44.

--- (2009a), "Sensory Qualities, Sensible Qualities, Sensational Qualities," in The Oxford Handbook to the Philosophy of Mind, eds. A. Beckermann, B. P. McLaughlin and S. Walter, Oxford: Oxford University Press, pp. 268-280.

--- (2009b), "Experience and Content," Philosophical Quarterly 59, 236: 429-451.

--- (2010), "Inverted Qualia," in The Stanford Enclopedia of Philosophy, ed. E. N. Zalta, accessed October 2011, URL = <http://plato.stanford.edu/entries/qualia-inverted/>.

Byrne, A. and D. R. Hilbert (2003), "Color Realism and Color Science," Behavioral and Brain Sciences 26: 3-64.

--- (2007), "Color Primitivism," Erkenntnis 66, 73: 73-105.

Campbell, J., (1993/1997), "A Simple View of Color," in Reality, Representation, and Projection, eds. J. Haldane and C. Wright. Oxford: Oxford University Press; reprinted in Readings on Color Volume 1: The Philosophy of Color, eds. A. Byrne and D. R. Hilbert. Cambridge, MA: MIT Press, pp. 177-190.

--- (2002), Reference and Consciousness, Oxford: Oxford University Press.

Cappelen, H. and E. Lepore (2007), Language Turned on Itself: The Semantics and Pragmatics of Metalinguistic Discourse, Oxford: Oxford University Press.

--- (2012), "Quotation," in The Stanford Enclopedia of Philosophy, ed. E. N. Zalta, accessed October 2011, URL = <http://plato.stanford.edu/entries/quotation/>.

Carruthers, P. (2004), "Phenomenal Concepts and Higher-Order Experiences," Philosophy and Phenomenological Research 68, 2: 316--36.

--- (2011), The Opacity of Mind: An Integrative Theory of Self-Knowledge, Oxford: Oxford University Press.

Carruthers, P. and B. Veillet (2007), "The Phenomenal Concept Strategy," Journal of Consciousness Studies 14, 9-10: 212-236.

Chalmers, D. J. (1996), The Conscious Mind: In Search of a Fundamental Theory, Oxford: Oxford University Press.

--- (2003), "The Content and Epistemology of Phenomenal Belief," in Consciousness: New Philosophical Perspectives, eds. Q. Smith and A. Jokic, Oxford: Oxford University Press, pp. 220-72.

--- (2004), "The Representational Character of Experience," in The Future for Philosophy, ed. B. Leiter, Oxford: Oxford University Press, pp. 153-181; reprinted in Chalmers (2010), pp. 339-379.

--- (2006), "Perception and the Fall from Eden," in Perceptual Experience, eds. T. S. Gendler and J. Hawthorne, Oxford: Oxford University Press, pp. 49-125; reprinted in Chalmers (2010), pp. 381-454.

--- (2007), "Phenomenal Concepts and the Explanatory Gap," in Phenomenal Concepts and Phenomenal Knowledge: New Essays on Consciousness and Physicalism, eds. T. Alter and S. Walter, Oxford: Oxford University Press, pp. 167-194; reprinted in Chalmers (2010), pp. 305-336.

--- (2010), The Character of Consciousness, Oxford: Oxford University Press.

Chen, X., M. Gabitto, Y. Peng, N. J. P. Ryba and C. S. Zuker (2011), "A Gustotopic Map of Taste Qualities in the Mammalian Brain," Science 333, 6047: 1262-1266.

Chisholm, R. (1957), Perceiving: A Philosophical Study, Itaca, NY: Cornell University Press.

Chomsky, N. (1965), Aspects of the Theory of Syntax, Cambridge, MA: MIT Press.

Churchland, P. M. (1981), "Eliminative Materialism and the Propositional Attitudes," Journal of Philosophy 78, February: 67-90.

--- (2007a), "On the Reality (and Diversity) of Objective Colors: How Color-Qualia Space is a Map of Reflectance Profile Space,Philosophy of Science 74, 2: 119-149.

--- (2007b), "The Evolving Fortunes of Eliminative Materialism," Contemporary Debates in Philosophy of Mind, eds. B. McLaughlin and J. Cohen, Malden, MA: Blackwell, pp. 160-181.

Clark, A. (1985), "Spectrum Inversion and the Color Solid," Southern Journal of Philosophy 23: 431-43.

--- (1993), Sensory Qualities, Oxford: Clarendon Press.

--- (2000), A Theory of Sentience, Oxford: Oxford University Press.

--- (2004), "Sensing, Objects, and Awareness: Reply to Commentators," Philosophical Psychology 17, 4: 553-579.

Coates, P. (2007), The Metaphysics of Perception: Wilfrid Sellars, Perceptual Consciousness, and Critical Realism, New York: Routledge.

Crane, T. (1992), The Contents of Experience: Essays on Perception, Cambridge, UK: Cambridge University Press.

--- (2002), "Introspection, Intentionality and the Transparency of Experience," Philosophical Topics 28, 2: 49-67.

--- (2005), "Intentionality," in The Oxford Companion to Philosophy, new edition, ed. T. Honderich, Oxford: Oxford University Press, pp. 438-439.

--- (2009a), "Is Perception a Propositional Attitude?" Philosophical Quarterly 59, 236: 452-469.

--- (2009b), "Representationalism," in Oxford Handbook to the Philosophy of Mind, eds. A. Beckermann, B. P. McLaughlin and S. Walter, Oxford: Oxford University Press, pp. 474-493.

Cohen,J. and M. Matthen (eds.) (2010), Color Ontology and Color Science, Cambridge, MA: MIT Press.

Cummins, R. (1996), Representations, Targets, and Attitudes, Cambridge, MA: MIT Press.

Cutter, B. and M. Tye (2011), "Tracking Representationalism and the Painfulness of Pain," Philosophical Issues 21, 1: 90-109.

Davidson, D. (1979), "Quotation," in Inquiries Into Truth And Interpretation, Oxford: Oxford University Press, pp. 79-92.

--- (1986), "A Coherence Theory of Truth and Knowledge," in Truth and Interpretation: Perspectives on the Philosophy of Donald Davidson, ed. E. LePore, Oxford: Basil Blackwell, pp. 307-319.

Davies, M. (1987), "Tacit Knowledge and Semantic Theory: Can a Five per cent Difference Matter?" Mind 96, 384: 441-462.

De Brigard, F. and J. Prinz (2010), "Attention and Consciousness," Wiley Interdisciplinary Reviews 1, 1: 51-59.

Dennett, D. C. (1969), Content and Consciousness, London: Routledge and Kegan Paul.

--- (1988), "Quining Qualia," in Consciousness in Contemporary Science, eds. A. J. Marcel and E. Bisiach, Oxford: Oxford University Press; reprinted in Philosophy of Mind: Classical and Contemporary Readings, ed. D. J. Chalmers, Oxford: Oxford University Press, pp. 226-246.

--- (1990), "The Myth of Original Intentionality," in Modeling the Mind, eds. K. A. Mohyeldin Said, W. H. Newton-Smith, R. Viale, and K. V. Wilkes, Oxford: Oxford University Press, pp. 43-62.

--- (1991), Consciousness Explained, Boston: Little, Brown and Co.

Descartes, R. (2006), Meditations, Objections, and Replies, eds. and tr. by R. Ariew and D. Cress, Indianapolis: Hackett Publishing Company.

Devitt, M. (1974), "Singular Terms," The Journal of Philosophy 71, 7: 183-205.

--- (1981), Designation, New York: Columbia University Press.

--- (2004), "The Case for Referential Descriptions," in Descriptions and Beyond, eds. M. Reimer and A. Bezuidenhout, Oxford: Oxford University Press, pp. 280-306.

Devitt, M. and K. Sterelny (1999), Language and Reality, 2nd Edition, Cambridge, MA: MIT Press.

Dretske, F. (1981), Knowledge and the Flow of Information, Cambridge, MA: MIT Press.

--- (1988), Explaining Behavior: Reasons in a World of Causes, Cambridge, MA: MIT Press.

--- (1995), Naturalizing the Mind, Cambridge, MA: MIT Press.

--- (1999), "The Mind's Awareness of Itself," Philosophical Studies 95: 103-124.

--- (2003), "Experience as Representation," Philosophical Issues 13: 67-82.

Egan, A. (2006), "Appearance Properties?" Noûs 40, 3: 495-521.

Evans, G. (1982), The Varieties of Reference, Oxford: Clarendon.

Fodor, J. A. (1975), The Language of Thought, Cambridge, MA: Harvard University Press.

--- (1983), The Modularity of Mind, Cambridge, MA: MIT Press.

--- (1984), "Semantics, Wisconsin Style," Synthese 59: 231-50.

--- (1987), Psychosemantics, Cambridge, MA: MIT Press.

--- (1990), A Theory of Content and Other Essays, Cambridge, MA: MIT Press.

--- (1998a), Concepts: Where Cognitive Science Went Wrong, Oxford: Clarendon Press.

--- (1998b), "There are No Recognitional Concepts, Not Even RED," Philosophical Issues 9, 1-14.

--- (2007), "The Revenge of the Given," in Contemporary Debates in Philosophy of Mind, eds. B. McLaughlin and J. Cohen, Malden, MA: Blackwell, pp. 105-116.

Fodor, J. and E. Lepore (1992), Holism: A Shopper's Guide, Cambridge, MA: Blackwell.

Forbes, G. (2006), Attitude Problems: An Essay in Linguistic Intensionality, Oxford: Oxford University Press.

Ford, J. (2011), "Tye-dyed Teleology and the Inverted Spectrum," Philosophical Studies 156, 2: 267-281.

Frege, G. (1892/1980), "On Sense and Reference," in Translations from the Philosophical Writings of Gottlob Frege, 3rd edition, eds. and trans. P. Geach and M. Black, Oxford: Blackwell.

Gärdenfors, P. (2004), "Conceptual Spaces as a Framework for Knowledge Representation," Mind and Matter 2, 2: 9-27.

Gendler, T. S. & J. Hawthorne (eds.) (2002), Conceivability and Possibility, Oxford: Oxford University Press.

Gibson, J. J. (1979), The Ecological Approach to Visual Perception, Boston: Houghton Mifflin.

Gilchrist, A., and A. Radonjić (2009), "Anchoring of Lightness Values by Relative Luminance and Relative Area," Journal of Vision 9, 9, 13: 1-10.

Glüer, K. (2009), "In Defence of a Doxastic Account of Experience," Mind and Language 24, 3: 297-327.

Goldstone, R. L., & A. T. Hendrickson, (2010), "Categorical Perception," Interdisciplinary Reviews: Cognitive Science, 1: 65-78.

Goodman, N. (1976), Languages of Art: An Approach to a Theory of Symbols, 2nd edition, Indianapolis: Hackett Publishing Company.

--- (1977), The Structure of Appearance, 3rd edition, Dordrecht-Holland: D. Reidel Publishing Company.

Graham, C. H., and Y. Hsia (1958), "The Discriminations of a Normal and a Color Blind Eye in the Same Person," Proceedings of the American Philosophical Society 102: 168-173.

Gunther, Y. (ed.) (2003), Essays on Nonconceptual Content, Cambridge, MA: MIT Press.

Hanna, R. and M. Chadha (2011), "Non-Conceptualism and the Problem of Perceptual Self-Knowledge," European Journal of Philosophy 19, 2: 184-223.

Hardin, C. L. (1993), Color for Philosophers: Unweaving the Rainbow, expanded edition, Indianapolis: Hackett.

--- (1997), "Reinverting the Spectrum," in Readings on Color, Vol. 1: The Philosophy of Color, eds. A. Byrne and D. Hilbert, Cambridge, MA: MIT Press, pp. 289-301.

Harman, G. (1987), "(Non-solipsistic) Conceptual Role Semantics," in New Directions in Semantics, ed. E. Lepore, London: Academic Press, pp. 55-81.

--- (1990), "The Intrinsic Quality of Experience," Philosophical Perspectives 4: 31-52.

--- (1996a), "Explaining Objective Color in Terms of Subjective Reactions," Philosophical Issues 7: 1-17.

--- (1996b), "Qualia and Color Concepts," Philosophical Issues 7: 75-79.

Harnad, S. (1987), Categorical Perception: The Groundwork of Cognition, Cambridge, UK: Cambridge University Press.

Haugeland, J. (1998), Having Thought, Cambridge, MA: Harvard University Press.

Heck, R. G. (2000), "Nonconceptual Content and the 'Space of Reasons'," The Philosophical Review 109, 4: 483-523.

Heed, T. (2010), "Touch Perception: How We Know Where We Are Touched," Current Biology 20, 14: R604-R606.

Held, R., Y. Ostrovsky, B. de Gelder, T. Gandhi, S. Ganesh, U. Mathur, and P. Sinha (2011), "The Newly Sighted Fail to Match Seen Shape with Felt," Nature Neuroscience 14: 551-553.

Hilbert, D. R. and M. E. Kalderon (2000), "Color and the Inverted Spectrum," in Color Perception: Philosophical, Psychological, Artistic and Computational Perspectives, ed. S. Davis, Oxford: Oxford University Press, pp. 187-214.

Hill, C. S. "The Paradox of Pain," in Pain: New Essays on Its Nature and the Methodology of Its Study, ed. M. Aydede, Cambridge, MA: MIT Press.

--- (2009), Consciousness, Cambridge, UK: Cambridge University Press.

Hill, C. S. and B. P. McLaughlin (1999), "There Are Fewer Things in Reality Than Are Dreamt of in Chalmers' Philosophy," Philosophy and Phenomenological Research 59, 2: 445-54.

Hopp, W. (2011), Perception and Knowledge: A Phenomenological Account, Cambridge, UK: Cambridge University Press.

Horgan, T. E. (1998), "Recognitional Concepts and the Compositionality of Concept Possession," Philosophical Issues 9: 27-33.

Horgan, T. and G. Graham (forthcoming), "Phenomenal Intentionality and Content Determinacy," in Prospects for Meaning, ed. R. Shantz, Amsterdam: DeGruyter, pp. 321-344.

Horgan, T. E. and J. L. Tienson (2002), "The Intentionality of Phenomenology and the Phenomenology of Intentionality," in Philosophy of Mind: Classical and Contemporary Readings, ed. D. J. Chalmers, Oxford: Oxford University Press, pp. 520-533.

Howard, J. D., J. Plailly, M. Grueschow, J.-D. Haynes and J. A. Gottfried (2009), "Odor Quality Coding and Categorization in Human Posterior Piriform Cortex," Nature Neuroscience 12, 7: 932-938.

Hume, D. (1739/2000), A Treatise of Human Nature, eds. D. F. Norton and M. J. Norton, Oxford: Oxford University Press.

Hurlburt, R. T. and E. Schwitzgebel (2007), Describing Inner Experience? Proponent Meets Skeptic, Cambridge, MA: MIT Press.

Hurley, S. R. (1998), Consciousness in Action. Cambridge, MA: Harvard University Press.

Husserl, E. (1970), Logical Investigations, 2 volumes, trans. J. N. Findlay, London: Routledge and Kegan Paul.

Jackson, F. (1977), Perception: A Representative Theory, Cambridge, UK: Cambridge University Press.

--- (1982), "Epiphenomenal Qualia," Philosophical Quarterly 32: 127-136.

--- (2007), "The Knowledge Argument, Diaphanousness, Representationalism," in Phenomenal Concepts and Phenomenal Knowledge: New Essays on Consciousness and Physicalism, eds. T. Alter and S. Walter, Oxford: Oxford University Press, pp. 52-64.

Johnston, M., (1992/1997), "How to Speak of the Colors," Philosophical Studies 68: 221-263; reprinted in Readings on Color Volume 1: The Philosophy of Color, eds. A. Byrne and D. R. Hilbert. Cambridge, MA: MIT Press, pp. 137-176.

Kanai, R. and N. Tsuchiya (2012), "Qualia," Current Biology 22, 10: R392-R396.

Kant, I. (1787/1998), Critique of Pure Reason, ed. and trans. P. Guyer and A. W. Wood, Cambridge, UK: Cambridge University Press.

Kaplan, D. (1989), ''Demonstratives', in Themes from Kaplan, eds. J. Almog, J. Perry, and H. Wettstein, Oxford: Oxford University Press, pp. 481--564.

Kind, A. (2003), "What's So Transparent about Transparency?" Philosophical Studies 115, 3: 225-244.

Kirk, R. (1974), "Sentience and Behaviour," Mind 83: 43-60.

Koch, C. and N. Tsuchiya (2007), "Attention and Consciousness: Two Distinct Brain Processes," Trends in Cognitive Sciences 11, 1: 16-22.

Kosslyn, S. M., W. L. Thompson, and G. Ganis (2006), The Case for Mental Imagery, Oxford: Oxford University Press.

Kotake, Y., H. Morimoto, Y. Okazaki, I. Fujita, and H. Tamura (2009), "Organization of Color-Selective Neurons in Macaque Visual Area V4," Journal of Neurophysiology 102: 15-27.

Kouider, S. and S. Dehaene (2007), "Levels of Processing During Non-Conscious Perception: A Critical Review of Visual Masking," Philosophical Transactions of the Royal Society B 362: 857-875.

Kriegel, U. (2006), "Theories of Consciousness," Philosophy Compass 1, 1: 58-64.

--- (2009), Subjective Consciousness: A Self-Representational Theory, Oxford: Oxford University Press.

Kripke, S. (1971), "Identity and Necessity," in Identity and Individuation, ed. M. K. Munitz, New York: New York University Press, pp. 135-164.

--- (1980), Naming and Necessity, Cambridge: Harvard University Press.

Kuehni, R. G. (2010), "Color Spaces and Color Order Systems: A Primer," in Color Ontology and Color Science, eds. J. Cohen and M. Matthen, Cambridge, MA: MIT Press, pp. 3-36.

Levine, J. (2001), Purple Haze: The Puzzle of Consciousness, Oxford: Oxford University Press.

--- (2003), "Experience and Representation," in Consciousness: New Philosophical Perspectives, eds. Q. Smith and A. Jokic, Oxford: Oxford University Press, 57-76.

--- (2010), "Phenomenal Experience: A Cartesian Theater Revival," Philosophical Issues 20, 1: 209-225.

Levin, J. (2000), "Dispositional Theories of Color and the Claims of Common Sense," Philosophical Studies 100: 151-174.

--- (2007), "What is a Phenomenal Concept?," in Phenomenal Concepts and Phenomenal Knowledge: New Essays on Consciousness and Physicalism, eds. T. Alter and S. Walter, Oxford: Oxford University Press, pp. 87-110.

Lewis, D. (1972), "Psychophysical and Theoretical Identifications," Australasian Journal of Philosophy 50, 3: 249-58.

--- (1986), "On the Plurality of Worlds," Oxford: Basil Blackwell.

Loar, B. (1990/1997), "Phenomenal States," in The Nature of Consciousness: Philosophical Debates, eds. N. Block, O. Flanagan and G. Guzeldere, Cambridge, MA: MIT Press, pp. 597-616.

--- (2003), "Transparent Experience and the Availability of Qualia," in Consciousness: New Philosophical Perspectives, eds. Q. Smith and A. Jokic, Oxford: Oxford University Press, pp. 77-96.

Locke, J. (1700/1975), An Essay Concerning Human Understanding, ed. P. H. Nidditch, Oxford: Clarendon Press.

Lycan, W. (1996a), Consciousness and Experience, Cambridge, MA: MIT Press.

--- (1996b), "Layered Perceptual Representation," Philosophical Issues 7: 81-100.

--- "Representational Theories of Consciousness," in The Stanford Enclopedia of Philosophy, ed. E. N. Zalta, accessed October 2011, URL = <http://plato.stanford.edu/entries/consciousness-representational/>.

--- (2006), "Pautz vs. Byrne & Tye on Externalist Intentionalism," Online Philosophy Conference, <https://webspace.utexas.edu/arp424/www/vs.pdf>.

Mandik, P. (2012), "Color-Consciousness Conceptualism," Consciousness and Cognition 21, 2: 617-631.

Martin, M. G. F. (1994) "Perceptual Content," in A Companion to the Philosophy of Mind, ed. S. Guttenplan, Oxford: Blackwell, pp. 463-471.

Mather, G., F. Verstraten and S. Anstis (1998), The Motion Aftereffect: A Modern Perspective, Cambridge, MA: MIT Press.

Matthen, M. (1988), "Biological Functions and Perceptual Content," Journal of Philosophy 85, 1: 5-27.

McDowell, J. (1994), Mind and World, Cambridge, MA: Harvard University Press.

--- (1998), "Having the World in View: Sellars, Kant, and Intentionality," Journal of Philosophy 95, 9: 431-491.

McGinn (1989), Mental Content, Oxford: Basil Blackwell.

--- (1991), The Problem of Consciousness, Oxford: Basil Blackwell.

--- (1996), The Character of Mind, 2nd edition, Oxford: Oxford University Press.

McNamara, T. (2005), Semantic Priming: Perspectives from Memory and Word Recognition, New York: Psychology Press.

Meehan, D. B. (2001), "Spatial Experience, Sensory Qualities, and the Visual Field," Proceedings of the Twenty-Third Annual Conference of the Cognitive Science Society, eds. J. D. Moore and K. Stenning, Mahwah, NJ: Lawrence Erlbaum Associates, Inc., pp. 623-627.

--- (2002), "Qualitative Character and Sensory Representation," Consciousness and Cognition 11, 4: 630-641.

--- (ms.), "The Qualiative Character of Spatial Perception," unpublished dissertation, CUNY Graduate Center.

Metzger, W. (1930), "Optische Untersuchungen am Ganzfeld," Psychologische Forschung 13: 6-29.

Millar, B. (2011), "Sensory Phenomenology and Perceptual Content," The Philosophical Quarterly 61, 244: 558-576.

Moore, G. E. (1903), "The Refutation of Idealism," Mind 12, 48: 433-453.

Morrison, John (2013), "Anti-Atomism about Color Representation," Noûs DOI: 10.1111/nous.12018.

Nagel, T. (1974), "What is it Like to be a Bat?," Philosophical Review 83, 4: 435-50.

Nanay, B. (forthcoming), "The Representationalism Versus Relationalism Debate: Explanatory Contextualism about Perception," European Journal of Philosophy.

Nida-Rümelin, M. (2007), "Transparency of Experience and the Perceptual Model of Phenomenal Awareness," Philosophical Perspectives 21, 1: 429-455.

Noë, A. (2004), Action in Perception, Cambridge, MA: MIT Press.

O'Regan, J. K. and A. Noë (2001), "A Sensorimotor Account of Vision and Visual Consciousness," Behavioral and Brain Sciences 24, 5: 939-1031.

Pace, M. (2005), "Review of Alva Noë's Action in Perception," Notre Dame Philosophical Reviews, published online: <http://ndpr.nd.edu/news/24903/?id=4761>.

Palmer, S. E. (1999), Vision Science: Photons to Phenomenology, Cambridge, MA: MIT Press.

Pautz, A. (2006), "Sensory Awareness is not a Wide Physicalism Relation," Noûs 40, 2: 205-240.

--- (2010a), "Why Explain Visual Experience in Terms of Content?" in Perceiving the World, ed. B. Nanay, Oxford: Oxford University Press, pp. 254-309.

--- (2010b), "Do Theories of Consciousness Rest on a Mistake?" Philosophical Issues 20, 1: 333-367.

--- (forthcoming), "The Real Trouble for Phenomenal Externalists: New Empirical Evidence for a Brain-Based Theory of Consciousness" in Consciousness Inside and Out: Phenomenology, Neuroscience, and the Nature of Experience, ed. Richard Brown, Springer, pp. xxx.

--- (ms.), "Why Consciousness Cannot Just be in the Head: A New Argument against Biological Theories," Second Annual Online Consciousness Conference, available online: <http://consciousnessonline.files.wordpress.com/ 2010/01/pautz-final-biological-2.pdf>.

Papineau, D. (1984), "Representation and Explanation," Philosophy of Science, 51: 550–72.

--- (2002), Thinking about Consciousness, Oxford: Oxford University Press.

--- (2007), "Phenomenal Concepts and the Materialist Constraint," in Phenomenal Concepts and Phenomenal Knowledge: New Essays on Consciousness and Physicalism, eds. T. Alter and S. Walter, Oxford: Oxford University Press, pp. 145-166.

Peacocke, C. (1983), Sense and Content: Experience, Thought, and their Relations, Oxford: Clarendon Press.

--- (1992), A Study of Concepts, Cambridge, MA: MIT Press.

--- (2001), "Does Perception Have a Nonconceptual Content?" The Journal of Philosophy 98, 5: 239-264.

--- (2008), "Sensational Properties: Theses to Accept and Theses to Reject," Revue Internationale de Philosophie 62: 7-24.

Perry, J. (2001), Knowledge, Possibility, and Consciousness, Cambridge, MA: MIT Press.

Persuh, M. and T. Ro (2012), "Context-Dependent Brightness Priming Occurs without Visual Awareness," Consciousness and Cognition 21, 1: 177-185.

Pitt, D. (2008), "Mental Representation," in The Stanford Enclopedia of Philosophy, ed. E. N. Zalta, accessed October 2011, URL = <http://plato.stanford.edu/entries/mental-representation/>.

Prinz, J. J. (2002), Furnishing the Mind: Concepts and Their Perceptual Basis, Cambridge, MA: MIT Press.

--- (2006), "Putting the Brakes on Enactive Perception," Psyche 12, 1: 1-19.

--- (2007), "Mental Pointing: Phenomenal Knowledge Without Concepts," Journal of Consciousness Studies 14, 9-10: 184-211.

--- (2012), The Conscious Brain: How Attention Engenders Experience, Oxford: Oxford University Press.

Quine, W.V. (1940), Mathematical Logic, Cambridge, MA: Harvard University Press.

--- (1960), Word and Object, Cambridge, MA: MIT Press.

--- (1961), "Reference and Modality," in From a Logical Point Of View, 2nd Edition, Cambridge, MA: Harvard University Press, pp. 139-159.

--- (1969), "Epistemology Naturalized," in Ontological Relativity and Other Essays, New York: Columbia University Press, pp. 69-90.

Raffman, D. (1995), "On the Persistence of Phenomenology," in Conscious Experience, ed. Thomas Metzinger, Exeter, UK: Imprint Academic, pp. 293-308.

Recanati, F. (2002), "The Fodorian Fallacy," Analysis 62, 4: 285-89.

Reid, T. (1764/1997), An Inquiry into the Human Mind on the Principles of Common Sense, ed. D. R. Brookes, University Park, PA: The Pennsylvania State University Press.

--- (1785/1969), Essays on the Intellectual Powers of Man, ed. B. A. Brody, Cambridge, MA: MIT Press.

Reimer, M. (1996), "Quotation Marks: Demonstratives or Demonstrations?" Analysis 56: 131-42.

Ro, T., N. S. Singhal, B. G. Breitmeyer and J. O. Garcia (2009), "Unconscious Processing of Color and Form in Metacontrast Masking," Attention, Perception, and Psychophysics 71, 1: 95--103.

Rosenthal, D. M. (1984), "Armstrong's Causal Theory of Mind," in Profiles: D. M. Armstrong, ed. R. J. Bogdan, Dordrecht: D. Reidel Publishing Company, pp. 79-120.

--- (1999), "The Colors and Shapes of Visual Experiences," in Consciousness and Intentionality: Models and Modalities of Attribution, ed. D. Fisette, Dordrecht: Kluwer Academic Publishers, pp. 95-118.

--- (2004), "Varieties of Higher-Order Theory," in Higher-Order Theories of Consciousness, ed. R. Gennaro, Philadelphia: John Benjamins.

--- (2005), Consciousness and Mind, Oxford: Clarendon Press.

--- (2010), "How to Think About Mental Qualities," Philosophical Issues 20, 1: 368-393.

--- (2012), “Awareness and Identification of Self,” in Consciousness and the Self: New Essays, eds. J. Liu and J. Perry, Cambridge,UK: Cambridge University Press, pp. 22-50.

Russell, B. (1903), The Principles of Mathematics, Cambridge, UK: Cambridge University Press.

--- (1912/1997), The Problems of Philosophy, Oxford: Oxford University Press.

--- (1948), An Inquiry into Meaning and Truth, London: George Allen and Unwin.

Schellenberg, S. (2011), “Perceptual Content Defended,” Noûs 45, 4: 714-750.

Schwitzgebel, E., “The Unreliability of Naïve Introspection,” Philosophical Review 117: 245-273.

--- (2010), “Introspection,” in The Stanford Enclopedia of Philosophy, ed. E. N. Zalta, accessed October 2011, URL = < http://plato.stanford.edu/entries/introspection/>.

Schwenkler, J. (2012) “On the Matching of Seen and Felt Shape by Newly Sighted Subjects,” i-Perception 3: 186-188.

Searle, J. R. (1983), Intentionality: An Essay in the Philosophy of Mind, Cambridge, UK: Cambridge University Press.

--- (1992), The Rediscovery of the Mind, Cambridge, MA: MIT Press.

Sellars, W. (1963), Science, Perception, and Reality, London: Routledge and Kegan Paul; reissued in 1991 by Ridgeview Publishing Company (Atascadero, CA).

--- (1967), Science and Metaphysics: Variations on Kantian Themes, London: Routledge and Kegan Paul; reissued in 1992 by Ridgeview Publishing Company (Atascadero, CA).

--- (1975), “The Structure of Knowledge Lecture I: Perception,” in Action, Knowledge and Reality: Studies in Honor of Wilfrid Sellars, ed. H.-N. Castañeda, Indianapolis: Bobbs-Merrill, pp. 295-347.

--- (1977), Philosophical Perspectives: Metaphysics and Epistemology, Atascadero, CA: Ridgeview Publishing Company.

Sextus Empiricus (2005), Against the Logicians, ed. and trans. R. Bett, Cambridge, UK: Cambridge University Press.

Shoemaker, S. (1975), “Functionalism and Qualia,” Philosophical Studies 27, 5: 291-315.

--- (1982), “The Inverted Spectrum,” Journal of Philosophy 79: 357-381.

--- (1990), “Qualities and Qualia: What’s in the Mind?” Philosophy and Phenomenological Research 50, Supplement: 109-31.

--- (1994a), “Self-Knowledge and “Inner Sense”: Lectures I-III,” Philosophy and Phenomenological Research 54, 2: 249-314.

--- (1994b), “Phenomenal Character,” Noûs 28, 1: 21-38.

--- (2000), “Phenomenal Character Revisited,” Philosophy and Phenomenological Research 60, 2: 465-467.

--- (2001), “Introspection and Phenomenal Character,” Philosophical Topics 28, 2: 247-73.

--- (2003), “Content, Character, and Color,” Philosophical Issues 13, 1: 253-78.

--- (2006), “On the Way Things Appear,” in Perceptual Experience, ed. T. S. Gendler and J. Hawthorne, Oxford: Oxford University Press, pp. 461-480.

--- (2007), “Thau on Perception,” Philosophical Studies 132, 3: 595-605.

Siegel, S. (2006), “Which Properties are Represented in Perception?” in Perceptual Experience, eds. T. S. Gendler and J. Hawthorne, Oxford: Oxford University Press, pp. 481-503.

--- (2010a), The Contents of Visual Experiences, Oxford: Oxford University Press.

--- (2010b), “The Contents of Perception,” in The Stanford Enclopedia of Philosophy, ed. E. N. Zalta, accessed October 2011, URL = <http://plato.stanford.edu/entries/perception-contents/>.

Siewert, C. P. (1998), The Significance of Consciousness, Princeton: Princeton University Press.

Smithies, D. and D. Stoljar (2012), Instrospection and Consciousness, Oxford: Oxford University Press.

Speaks, J. (2005), “Is There a Problem about Nonconceptual Content?” Philosophical Review 114, 3: 359-398.

--- (2009), “Transparency, Representationalism, and the Nature of Perceptual Content,” Philosophy and Phenomenological Research 79, 3: 539-573.

--- (2011), “Spectrum Inversion without a Difference in Representation is Impossible,” Philosophical Studies 156, 3: 339-361.

Spillman, L., W. Laskowski, and K. W. Lange (2000). “Stroke-blind for Colors, Faces, and Locations: Partial Recovery after Three Years,” Restorative Neurology and Neuroscience17: 89-103.

Stampe, D. (1977), "Towards a Causal Theory of Linguistic Representation," Midwest Studies in Philosophy 2: 42-63.

Stalnaker, R. (1987), Inquiry, Cambridge, MA: MIT Press.

Sterelyny, K. (1990), The Representational Theory of Mind, Oxford: Basil Blackwell.

Stich, S. and T. Warfield (eds.) (1994), Mental Representation, Oxford: Basil Blackwell.

Stoljar, D. (2005), "Physicalism and Phenomenal Concepts," Mind and Language 20, 469-94.

Tarski, A. (1933), "The Concept of Truth in Formalized Languages," reprinted in Tarski (1983), Logic, Semantics, Metamathematics, 2nd Edition, Indianapolis: Hackett, pp. 152-278.

Thau, M. (2002), Consciousness and Cognition, Oxford: Oxford University Press.

Thompson, B. (2007), "Shoemaker on Phenomenal Content," Philosophical Studies 135: 307-334.

--- (2009), "Senses for Senses," Australasian Journal of Philosophy 87, 1: 99-117.

Treisman, A. (1998), "Feature Binding, Attention and Object Perception," Philosophical Transactions of the Royal Society of London B 353: 1295-1306.

Tye, M. (1995), Ten Problems of Consciousness: A Representational Theory of the Phenomenal Mind, Cambridge, MA: MIT Press.

--- (1996), "Perceptual Experience Is a Many-Layered Thing," Philosophical Issues 7: 117-126.

--- (2000), Consciousness, Color, and Content, Cambridge, MA: MIT Press.

--- (2002), "Visual Qualia and Visual Content Revisited," in Philosophy of Mind: Classical and Contemporary Readings, ed. D. J. Chalmers, Oxford: Oxford University Press, pp. 447-456.

--- (2003a), "Blurry Images, Double Vision, and Other Oddities: New Problems for Representationalism?" in Consciousness: New Philosophical Perspectives, eds. Q. Smith and A. Jokic, Oxford: Oxford University Press, pp. 7-32.

--- (2003b), "A Theory of Phenomenal Concepts," in Minds and Persons, ed. A. O'Hear, Cambridge, UK: Cambridge University Press, pp. 91-105.

--- (2006), "Nonconceptual Content, Richness, and Fineness of Grain," in Perceptual Experience, eds. T. S. Gendler and J. Hawthorne, Oxford: Oxford University Press, pp. 504-530.

--- (2007), "Qualia," in The Stanford Enclopedia of Philosophy, ed. E. N. Zalta, accessed October 2011, URL = <http://plato.stanford.edu/entries/qualia/>.

--- (2011), Consciousness Revisited: Materialism without Phenomenal Concepts, Cambridge, MA: MIT Press.

van Boxtel, J. J. A., N. Tsuchiya and C. Koch (2010), "Consciousness and Attention: On Sufficiency and Necessity," Frontiers in Psychology 1, 217: 1-13.

Vendler, Z. (1972), Res Cogitans: An Essay in Rational Psychology, Ithaca, NY: Cornell University Press.

Watkins, M. (2005), "Seeing Red: The Metaphysics of Colour without the Physics," Australasian Journal of Philosophy 83, 1: 33-52.

Weisberg, J. (ms.), The Problem of Consciousness: Mental Appearance and Mental Reality, unpublished dissertation, CUNY Graduate Center.

Winawer J., and N. Witthoft, M. Frank, L. Wu, A. Wade, and L. Boroditsky, (2007), "The Russian Blues Reveal Effects of Language on Color Discrimination," Proceedings of the National Academy of Science. 104, 19:7780-7785.

Zalta, E. (2001), "Fregean Senses, Modes of Presentation, and Concepts," Philosophical Perspectives 15: 335-359.