

PHENOMENAL CONCEPTS

BY DOUGLAS PARVIN

A dissertation submitted to the
Graduate School—New Brunswick
Rutgers, The State University of New Jersey
in partial fulfillment of the requirements

for the degree of

Doctor of Philosophy

Graduate Program in Philosophy

Written under the direction of

Brian McLaughlin

and approved by

New Brunswick, New Jersey

May, 2009

ABSTRACT OF THE DISSERTATION

Phenomenal Concepts

by Douglas Parvin

Dissertation Director: Brian McLaughlin

I explore various claims about the nature of phenomenal concepts and isolate two recurring intuitions. The first involves the epistemological role of phenomenal concepts: a phenomenal concept is supposed to be a concept of a type of experience that must be possessed by a subject who knows what it is like to have an experience of the type in question. The second involves the importance of experience: a phenomenal concept is supposed to be a concept of a type of experience that can be possessed only by a subject who has had an experience of the type in question. Most accounts of phenomenal concepts have stipulations designed to satisfy both these conditions. I argue, however, that they cannot jointly be satisfied.

We thus face a choice: either we can possess phenomenal concepts of types of experiences we haven't had, or a phenomenal concept is not required for phenomenal knowledge. I argue that the latter is unacceptable, as the idea of a phenomenal concept is inextricably tied to issues involving the relationship between phenomenal knowledge and non-phenomenal knowledge.

I defend a recognitional account of phenomenal concepts, whereby a subject possesses a phenomenal concept partly in virtue of being able to recognize an experience as being of a certain type and which does not require having had an experience of the type in question. I consider and reject the rival “quotational” account, which holds that a phenomenal concept actually contains its referent as a proper part.

The latter part of my dissertation is an analysis of some prominent antiphysicist arguments through the lens of phenomenal concepts. I consider, especially, a theme that runs through them, which is what Brian Loar calls ‘the semantic premise’, and which Stephen White has recently argued for: the claim that any true identity statement that involves noncontingent modes of presentation on both sides of the identity must be a priori.

Acknowledgements

Much thanks to Brian McLaughlin, who guided my efforts with infinite patience. Thanks also to Barry Loewer, Dean Zimmerman, Jerry Fodor, David Chalmers, and Daniel Stoljar, for their illuminating comments. Finally, thanks to the many graduate students at Rutgers University who provided ample suggestions and criticisms.

Dedication

For Mom, Dad, Beth, and Jenny.

Table of Contents

Abstract	ii
Acknowledgements	iv
Dedication	v
1. Phenomenal Concepts and the Phenomenal Concept Strategy	1
1.1. Introduction	1
1.2. Chapter Overview	4
1.3. The Knowledge Argument and the Ability Hypothesis	8
1.3.1. The Knowledge Argument	8
1.3.2. The Ability Hypothesis	10
1.4. The Explanatory Gap	16
1.4.1. The Nature and Significance of the Explanatory Gap	17
1.4.2. The Explanatory Gap and Phenomenal Concepts: Reply to Chalmers	21
2. The Demonstrative-Recognitional Account of Phenomenal Concepts	25
2.1. Type-Demonstrative Phenomenal Concepts	25
2.1.1. Thick and Thin Demonstrative Phenomenal Concepts	29
2.2. Do We Possess Enough Introspectively Applied Type-demonstrative Concepts to Account for the Extent of Our Phenomenal Knowledge?	33

3. Acquaintance, Revelation, and the Quotational Account of Phenomenal Concepts	39
3.1. Chapter Introduction	39
3.2. The Challenge of Blindsight	41
3.3. Quotational Accounts of Phenomenal Concepts	44
3.4. Awareness and Acquaintance	49
3.4.1. An Innocent Notion of Awareness	49
3.4.2. The Revelation Intuition	52
3.4.3. There are Numerous Ways to Acquire Phenomenal Concepts	54
3.5. Postscript: Imaginative Abilities and Phenomenal Representation	57
4. Phenomenal Empiricism	62
4.1. The Thesis of Phenomenal Empiricism	62
4.2. The Leading Theories Entail Phenomenal Empiricism	64
4.3. Why Phenomenal Empiricism is False	65
4.3.1. Counterexamples	65
4.3.2. The Semantics of Phenomenal Concepts	71
4.4. Purely Recognitional Phenomenal Concepts Vs. Phenomenal Demonstratives	72
5. Recognitional Concepts: Reply to Fodor	77
5.1. Fodor's Argument	78
5.2. Nominalism, Realism, and Fodor's Argument	81
5.2.1. Nominalist and Realist Conceptions of Recognitional Concepts	81
5.2.2. Fodor's Assumption of a Nominalist Conception of Recognitional Concepts	82

5.2.3. Compositionality on a Realist Conception of Recognitional Concepts	83
5.3. Are There Complex Phenomenal Concepts?	89
6. The Phenomenal Concept Strategy: Reply to Stoljar	93
6.1. Chapter Introduction	93
6.2. Varieties of the Phenomenal Concept Strategy	94
6.3. Stoljar’s First Argument	98
6.4. Stoljar’s Second Argument	101
7. The Property Dualism Argument: Reply to White	105
7.1. Introduction to the Argument	105
7.1.1. The Classic Property Dualism Argument	106
7.1.2. The Phenomenal Concept Reply	108
7.2. The New Property Dualism Argument	109
7.2.1. Representational Coherence	110
7.2.2. The “A Priori Condition”	112
7.2.3. Recognitional Concepts vs. (Mere) Recognitional Abilities	116
7.3. The Importance of Type-Identity to A Posteriori Physicalism . . .	120
7.4. An Additional Problem for Type-Demonstrative Accounts of Phenomenal Concepts	128
7.4.1. Phenomenal Concepts, Demonstratives, and Pure Indexicals	129
References	133

Chapter 1

Phenomenal Concepts and the Phenomenal Concept Strategy

1.1 Introduction

This work will serve largely as a defense of the phenomenal concept strategy. I will lay out what I believe to be the strongest physicalist account of phenomenal concepts, defend it from some objections, and apply the account to some recent formulations of antiphysicalist arguments.

In recent years, the notion of a phenomenal concept has figured centrally in the debate over the place of phenomenal consciousness in nature. A phenomenal concept is supposed to be a certain kind of concept that plays at least two roles. First, a phenomenal concept has the role of denoting a phenomenal property.¹ The second role is epistemic. It is generally accepted that a phenomenal concept of a phenomenal property must be possessed by a subject who knows what it is like to have the phenomenal property in question.² I assume in this paper the prevailing view that knowing what it is like to have a certain experience requires knowing certain facts, and so requires possessing concepts sufficient to represent

¹This work will not make any substantive assumptions about the general nature or existence of properties. The expressions 'phenomenal property' and 'type of experience' should, for present purposes, be understood to be synonymous. Thus, experiences are understood to be typed by their subjective feels, that is, by their what-it-is-like for the subject aspect. All that is required for now is that we possess concepts which intelligibly cut the ocean of phenomenal experience and pick out parts of it in a systematic way.

²The centrality of 'knowing what it is like' to discussions of phenomenal consciousness comes largely from (Nagel, 1974).

those facts (Tye (2001), Chalmers (2003), and Loar (1997)).

It is assumed both by leading dualists and by leading physicalists that understanding the nature of phenomenal concepts is crucial for assessing many well-known antiphysicalist arguments, including the Knowledge Argument (Jackson, 1982, 1986), the Conceivability Argument (Chalmers, 1996), and the Property Dualism Argument, or “Distinct Properties Argument”, as it’s sometimes called (White, 2006). Physicalists employ a certain strategy that has come to be called “the phenomenal concept strategy”, following (Stoljar, 2005). The strategy involves an appeal to the special nature of phenomenal concepts against these dualist arguments, which purport to show that various physicalist claims (such as psychophysical identity claims) are a priori false.³ Dualists counter by disputing the physicalist accounts of phenomenal concepts, offering their own accounts, and using them to reformulate and sharpen the arguments in question.

The central move in the phenomenal concept strategy is to grant a claim that dualists have made for some time: that phenomenal concepts are conceptually independent of physical and functional concepts, in the sense that there are no analytic entailments from physical or functional facts to phenomenal facts. But while dualists have taken this claim as a premise for numerous arguments against physicalism, the proponent of the phenomenal concept strategy argues that the truth of this claim actually supports physicalism, as it can serve to explain the existence of various explanatory gaps between phenomenal and physical truths. The proponent of the phenomenal concept strategy thus sides with the dualist, in this respect, against proponents of analytic (or a priori) physicalism, functionalism, or behaviorism.

³Throughout this work, I count functional properties that are always realized in our world by physical properties as themselves physical properties. Similarly, I count as physicalist those versions of functionalism that take it to be nomologically necessary that phenomenal properties are physically realized.

While dualists seek to explain these explanatory gaps as ontologically driven, the physicalist explains the gaps as expected outcomes of ontologically innocent contingent facts about our psychologies. Since we happen to conceptualize phenomenal properties in a way which does not connect in certain ways with how we conceptualize other physical, functional, and behavioral properties, such gaps are to be expected. But they are merely cognitive gaps, the physicalist argues, and have no ontological or evidentiary significance.

The success of the phenomenal concept strategy depends, in part, on whether conceivability entails possibility. If it does, the strategy is stillborn. So much of the debate over phenomenal concepts has been focused on this central question in modal epistemology that a casual observer might take the debate over phenomenal concepts simply to be an interesting application of the general question of the relationship between conceivability and possibility. Up until the past five or ten years, the casual observer would have been more or less right. Accounts of phenomenal concepts and criticisms of such accounts were analyzed largely through the lens of the conceivability-possibility debates.

As important as such debates are, however, this work will not focus on issues in modal epistemology (except for a bit in chapter six). Rather, it will focus on more recent developments in the ongoing debate over the nature of phenomenal concepts and the viability of the phenomenal concept strategy.

Over the last five or ten years, dualists and skeptics have opened up a new front against the phenomenal concept strategy. They have claimed that the physicalist accounts of phenomenal concepts themselves are plainly inadequate in various respects. If their criticisms are correct, then the phenomenal concept strategy fails even if the physicalist succeeds in refuting the claim that there is a robust link between conceivability and possibility. For if there is no physicalistically respectable account of phenomenal concepts to be had, then there will be a new

and different explanatory gap between the physical and conceptual. This would be just as problematic for the physicalist as an explanatory gap between the physical facts and facts regarding phenomenal concepts (Chalmers, 2006).

It is this new front on the phenomenal concept strategy that this work in large part addresses.

The account of phenomenal concepts I defend is one in which a phenomenal concepts is understood to be a special kind of introspectively applied recognitional concept. This sort of account is not new; among physicalists who write on the subject, it is perhaps the dominant view. With the exception of some modifications I propose in chapter four, my account is quite similar to the influential ‘demonstrative-recognitional’ account defended in Loar’s highly influential *Phenomenal States*(1990/97). Since its publication, there have been numerous criticisms of the account. These criticisms have been leveled not just by dualists, but by physicalists and agnostics as well, and this work serves partly as a defense of the account from these recent attacks. The other part is a refinement and emendation of the demonstrative-recognitional account of phenomenal concepts.

1.2 Chapter Overview

The remainder of this chapter sketches the knowledge argument and the argument from an explanatory gap, explores some physicalist responses other than the phenomenal concept strategy, and argues that these other responses leave something to be desired. I end the chapter with an application of the phenomenal concept strategy to a formulation of the explanatory gap argument recently put forward by Chalmers (2006).

In Chapter Two, I lay out some of the geography of the battle over the nature of phenomenal concepts, especially as it has arisen in response to the influential

type-demonstrative account of phenomenal concepts. I offer a limited defense of the view, paying special attention to an influential argument by Raffman (1995), who claims that our stock of type-demonstrative phenomenal concepts is insufficient in number to account for all the phenomenal properties we know what it's like to have.

In Chapter Three, I consider the charge that demonstrative concepts (and other related recognitional concepts) are too “thin” to account for the rich and substantive nature of the knowledge we typically gain upon acquaintance with a phenomenal property. This charge can be expressed in at least two ways. Some insist that phenomenal concepts require special modes of presentation or a special, irreducibly mental relation between the tokener of a phenomenal concept and the property the concept denotes. Others simply charge that demonstratives and other recognitional concepts fail to serve as phenomenal concepts because they do not capture some purportedly special, substantive aspect of phenomenal knowledge. At heart, these two charges converge on the following assertion: deploying a phenomenal concept must involve more than simply deploying an indexical or demonstrative.

I suggest that the charge is spurious, and that two influential proposals it motivates are misguided. The first proposal I consider is the proposal that we need to posit a basic relation of “acquaintance” between a subject and a phenomenal property. Such a relation, I argue, is either hopelessly murky or can be explicated as a certain unmysterious type of causal relation. The second proposal I reject is that phenomenal concepts in some sense “contain” or “quote” their referents (or tokens thereof). I suggest that, depending on how the account is spelled out, it is either implausible that we possess such concepts in many cases where we uncontroversially possess phenomenal knowledge, or that the quotational nature of the concept is dispensable window dressing that can be seen to be functionally

equivalent to a causal relation between a phenomenal concept and the property it denotes.

I end the chapter with an exploration of the relationship between phenomenal concept possession and imaginative abilities and consider the possibility of phenomenal intentionality, as suggested by Loar (2003), and its implication for a theory of phenomenal concepts.

In Chapter Four, I argue that the leading theories of phenomenal concepts, including the orthodox demonstrative account, are committed to an implausible doctrine I call *Phenomenal Empiricism*. Phenomenal Empiricism holds that knowing what it is like to have a phenomenal property requires having had the property in question. I give reasons to reject the doctrine, and a fortiori those theories that entail it. I present as an alternative a non-demonstrative recognitional account of phenomenal concepts and lay out a simple theory of the semantics of phenomenal concepts. This account, I argue, retains the virtues of the demonstrative account without committing to Phenomenal Empiricism.

Chapter Five considers an important challenge to the existence of recognitional concepts recently argued by Fodor (1998). Fodor rejects recognitional concepts on the grounds that they fail to compose. I suggest that Fodor's worries, while perhaps applicable to certain theories of recognitional concepts, do not apply to phenomenal concepts, for two reasons. First, I claim that Fodor's line of reasoning that leads him to conclude that recognitional concepts do not compose is sound only on a particular nominalist theory of recognitional concepts. I argue that phenomenal concepts should not be construed in this nominalist way, and that his argument therefore fails to show that there cannot be recognitional concepts of the sort that I take phenomenal recognitional concepts to be. Second, I argue there is no reason to suppose that there are any complex phenomenal concepts (though I grant that there certainly are other sorts of complex concepts with

phenomenal concepts as constituent parts). Once this is understood, Fodor's worries about compositionality go away.

Chapter Six considers a recent and influential attack on the phenomenal concept strategy by Stoljar (2005). Stoljar (the coiner of the phrase 'phenomenal concept strategy') grants that the strategy may show why subjects don't know what it's like to have phenomenal properties they've never had, and thus perhaps can explain why Jackson's Mary doesn't know what it's like to see red in her black-and-white room. He argues, however, that the strategy is powerless to explain why such knowledge is a posteriori, as the physicalists employing the strategy would have it.

I suggest that Stoljar misconstrues the way physicalists have used (or at least should use) the phenomenal concept strategy. The goal is not to support Phenomenal Empiricism nor to demonstrate that there are no a priori arguments for physicalism: rather it is to leave open the epistemic possibility that physicalism is true a posteriori (a posteriori physicalism will find its positive support elsewhere). By misconstruing the project, Stoljar attacks a straw man. Contingent facts about human concept acquisition have no bearing on the strategy (or at least should have no bearing on it). The point of the strategy is served by noting that a physical and a phenomenal concept might co-refer without the coreference being knowable a priori. Or so I argue.

Finally, Chapter Seven considers the property dualism argument, which, while more narrowly targeted than the knowledge and conceivability arguments, presents perhaps a greater challenge to identity theories. The reason for this is that the property dualism argument does not require possibility to follow from conceivability. The argument has enjoyed a resurgence of late thanks in large part to Steve White's (2006) careful elaboration and strengthening of the argument in response to the phenomenal concept strategy, as well as Ned Block's (2006) detailed

response to the argument. Block's response depends crucially on a quotational account of phenomenal concepts, but I suggest that a simpler recognitional account of phenomenal concepts can do the job for the physicalist.

I claim that considerations from the property dualism argument lend further support to a non-demonstrative account of recognitional phenomenal concepts, as they, and not demonstrative concepts, can uncontroversially be taken to refer without any associated mode of presentation. It is this fact, I suggest, which is key to defeating the property dualism argument.

The last section of Chapter Seven explores whether recognitional phenomenal concepts have associated modes of presentation. If they do, responding adequately to the property dualism argument becomes quite challenging. I suggest that there are good reasons to suppose that demonstrative recognitional concepts typically *do* have associated modes of presentation, but that non-demonstrative recognitional concepts, of the sort argued for in chapter four, clearly do not.

1.3 The Knowledge Argument and the Ability Hypothesis

1.3.1 The Knowledge Argument

Let us begin the discussion by treading some familiar ground, so we may get a sense of the sorts of questions to which the phenomenal concept strategy has been applied, of some other lines of defense physicalists have employed, and of some of the ways in which these alternative defenses have been seen by many to be inadequate.

Jackson (1982, 1986) presents a well-known thought experiment designed to show that there are facts about the subjective character of experience that no

physicalist theory can capture.⁴ He asks us to consider Mary, the color-challenged scientist:

Mary is confined to a black-and-white-room, is educated through lectures relayed on black-and-white television. In this way she learns everything there is to know about the physical nature of the world. She knows all the physical facts about us and our environment. [...] If physicalism is true, she knows all there is to know. [...] It seems however, that Mary does not know all there is to know. For when she is let out of the black-and-white room or given a color television, she will learn what it is like to see something red, say. This is rightly described as learning - she will not say “ho, hum.” Hence, physicalism is false.⁵

Schematically, we may present the argument as follows:

- (1) Mary (before her release) knows every truth about human experience derivable from physical facts.⁶
- (2) Mary, upon her release, learns what it’s like to see something red.
- (3) Mary (before her release) does not know every truth about human experience. (1, 2)
- (4) Therefore, there are truths about human experience that are not derivable from physical facts. (1, 3)

There has been much written about this argument, and so I will be brief in sketching the most prominent physicalist response to it, as set-up to an introduction of the phenomenal concept strategy as marking an improvement in the physicalist toolkit.

⁴Jackson no longer believes that the knowledge argument succeeds, holding that it is undermined by a form of representationalism he has come to endorse. Although Jackson no longer supports the argument, it is alive and well in the philosophical literature. I will not be addressing representationalism in this thesis.

⁵(Jackson, 1982), p. 292

⁶We may include biological, chemical, neurofunctional, and behavioral facts, as well, insofar as such facts are not derivable from basic physical facts.

Before getting into the ability hypothesis, however, it is worth noting another extant physicalist strategy. This strategy, which I will not dwell on, is to doggedly deny (2). Dennett (1991), for example, questions the reliability of any intuitions we might have regarding Mary’s pre-release knowledge or post-release surprise. He grants that *you or I* would learn something upon release, but we’re not super-scientists with such unearthly memory and mental computing power that we are able to store or instantly calculate infinitely many facts (remember, Mary knows all the physical facts, not just the basic physics formulas). Such an alien being, Dennett suggests, might figure out what it’s like to see red before release and may well, upon seeing a rose for the first time, say “ho, hum.”

This line has not proven particularly persuasive among philosophers, as it seems to neglect a core component of the intuition that Jackson elicits. It just doesn’t seem that *any* extra bit of physical information will put Mary *even slightly* on the path to knowing what it’s like to see something red. Nor, does it seem, will *any* extra bit of processing power. The problem seems not to have anything to do with the *quantity* of facts she has access to, nor with the power of her mind to process those facts, but rather with the *nature* of the facts themselves. They seem categorically ill-suited to the task at hand.

1.3.2 The Ability Hypothesis

A much more popular line against Jackson is what has come to be known as the “ability hypothesis,” advocated most notably by Lewis (1988) and Nemirow (1990). The ability hypothesis grants (2). That is, it grants that Mary learns something upon her release. However, it denies the move from (2) to (3). According to Lewis, such a move depends on an unargued implied premise, which he calls the hypothesis of phenomenal information:

Hypothesis of Phenomenal Information. *Besides physical information, there*

is an irreducibly different kind of information to be had: phenomenal information.

If the hypothesis of phenomenal information is true, then two worlds might be alike in every physical respect but differ phenomenally. This would contradict various supervenience theses commonly held by physicalists, such as the thesis that two worlds alike in all physical respects are alike simpliciter. So the challenge to the physicalist is to deny the hypothesis while granting that Mary learns something. This involves two parts. First, the physicalist must cast doubt on the plausibility of the hypothesis of phenomenal information. Second, the physicalist must offer a more plausible account of Mary's learning, which does not imply the hypothesis of phenomenal information.

Two of the worries Lewis raises for the hypothesis of phenomenal information don't strike me as very serious. First, Lewis notes that the intuition which grounds the hypothesis isn't merely the intuition that there is information that is unavailable from the sort of physics lessons Mary would be receiving within her black-and-white room. Rather, the intuition is that there is information that is unavailable from lessons of *any kind* that could be given to the inexperienced, including lessons from the textbook of a full-blown "qualia freak" (as Jackson likes to call himself). Hence, Lewis claims, the knowledge argument is self-refuting because it "proves too much".

It seems to me that the knowledge argument would be unsuccessful against such a dualistic theory because the first premise can't be realized. One way in which dualism differs from physicalism is that there is no generally held presumption that a complete theory of the world (which presumably has as part of its ontology irreducible phenomenal information) could ever be taught completely through books or over television, or in any objective or discursive manner. Even if one believed that such a theory would be couched, in part, in a heretofore undiscovered language of "objective phenomenology", there is no reason to

suppose that it could be taught in the standard ways we take it that physics or neuroscience could be taught. So the fact that such a dualist theory could not be learned completely from a textbook does not cast doubt on the theory, because there is no presumption that it could be learned in the ways we presume physics can be learned.

Lewis also finds it peculiar that the hypothesis of phenomenal information treats information as the elimination of possibilities. It is peculiar, he thinks, because when it comes to the putative phenomenal information it is not typically the case that we know in advance the possibilities that are candidates for elimination. When I learn the location of an object in space-time, for instance, there is a clear sense in which I have eliminated the other possibilities (i.e., the other points in space-time). But such does not seem to be the case when one, say, tastes vegemite for the first time (to borrow Lewis's pet example). In such a case, one would have no set of possible sensations to eliminate.

Before tasting, Lewis argues, one can't fully pose the question: is it *this* way or *that*? This strikes me as not at all obvious. For one thing, there are all sorts of ways I can imagine an experience being, and there are all sorts of ways that I can, after having the experience, qualitatively differentiate it from other experiences. The vegemite might have been sweeter, less salty, or perhaps even chocolaty, for all I know.

Granted, it is less clear what alternatives pre-release Mary could have in mind for red, but she may, for instance, make inter-modal comparisons. She may imagine, for example, that hue is like pitch, brightness is like loudness, and that saturation is like tonality. She may, in retrospect, say that some of her guesses were better than others. She would hold that some possibilities had been eliminated. Granted, the range of antecedently conceptualized possibilities might have been quite circumscribed. Still, even one eliminated possibility invalidates this

line of response.

Another aspect of the hypothesis of phenomenal information that Lewis finds peculiar is that it apparently leads in the direction of epiphenomenalism. For if phenomenal information is logically independent of physical information, then there are innumerable many logically possible worlds that are qualitatively identical in all physical respects, but which differ in some phenomenal aspect. Oddly enough, Mary's exclamations upon seeing red for the first time could be the same whatever the phenomenal information she is presented with. It is beyond the scope of this paper to critique epiphenomenalism in general, and to assess its epistemological implications. As a self-explanatory objection to the knowledge argument, however, it is unpersuasive. Though it strikes many as absurd, epiphenomenalism is alive and well as (perhaps) the most common dualist position endorsed by philosophers. So insofar as it is presented as an objection, it is question-begging in the present context.

Nemirow (1990) suggests that to know what an experience is like is to possess the ability to place one's self in a state representative of the experience in question. Having experiences, he suggests, puts one in position to remember and to imagine, and thus confers on a subject knowledge of what an experience is like. Seeing red (and knowing that it is red that is seen), for example, imparts the ability to make future attributions of red more easily, the ability to imagine related (or opposed) colors, and the ability to gain other, related abilities. Knowing what it's like, Nemirow claims, is a knowing-*how* as opposed to a knowing-*that*. It is not, according to this view, a matter of the acquisition of information. A fortiori, it is not a matter of the acquisition of phenomenal information.

If Nemirow is right, then it should come as no surprise that we often have trouble communicating phenomenal knowledge through language. Language is well-suited to information transfer, but it is often poorly suited to teaching an

ability. I may have trouble telling you how to visualize red, but this should not be surprising, on Nemirow's view. I might have as much trouble telling you how to wiggle your ears or lower your blood pressure — two clear cases of ontologically innocent knowing-how. I can't learn to tango (certainly not well, at least) by reading a book, and I would be a fool to read a book on karate and think that I've acquired the know-how to go down to the dojo and spar with a black-belt.

The black belt at the dojo knows karate; I do not. It may be the case that I can describe, in great detail, the mechanics of a round-house kick, while the black belt can only demonstrate it, being lost for words when it comes to offering an accurate description of the act. Whatever know-that I may possess, I would fall to the floor were I to try the kick. Knowing, in the case of karate, seems to have little if anything to do with information. For I memorized the book and have all the information. In Lewis's terminology, I can be said to have the information because I am able to eliminate alternative possibilities.

It would be a mistake, Lewis and Nemirow claim, to think that 'knowing' is ambiguous between 'knowing how' and 'knowing that'. It is better to think of 'knowing' as vague. There may be a continuum of cases between knowledge being a matter of pure information possession and it being a matter of pure ability possession. Knowing what it is like to have an experience, Lewis and Nemirow suppose, falls on the latter end of the spectrum.

As many have noted, it seems that there's often a sense in which knowing-how can be properly re-described as knowing-that. Let us return to our karate expert, who, by hypothesis, can execute the moves perfectly but is unable to describe what he is doing. We agree that he knows-how. Does he also know-that?

Suppose he sees one of his students clumsily executing a move and that he corrects the student by guiding him through the correct motion. It seems he has

knowledge-that — specifically, knowledge that the move looks or feels like such-and-such. In Lewis’s terms, he is able to eliminate the alternative possibilities. I, who have read the book, may have information the master lacks. But he has information I lack. He is able to pick out the correct kick from among the possibilities (he has seen many poorly executed kicks before, so nothing surprises him). I cannot do this. I almost certainly will not be able to do this if the kick is at full speed. If possessing information is having an ability to distinguish actualities from possibilities, then the karate expert possesses information that I lack.

Bringing this back to Jackson’s Mary, it seems that when she sees red she gains knowledge-that (i.e., knowledge that seeing red is like such-and-such). Unless we are willing to fortify the minimal physicalist picture with the much stronger claim that all experience-types must be characterized criterially through the language of physics, it seems possible that Mary could correct her book knowledge if something seemed red to her under optimal viewing conditions but was not counted as red by the physical theory.

Loar (1990/97) gives more reasons to reject the ability hypothesis, arguing that phenomenal knowledge involves “genuinely predicative components of judgment” that the ability hypothesis cannot account for. He gives, as an example, the thought “if coconuts did not have *this* taste, then Q”. To evaluate the thought, a subject must know what it is like to taste coconut. But in this case, there is no easy way to construe this phenomenal knowledge as possession of an ability.

Loar also notes that we entertain thoughts about the phenomenal states of other people, but that such thoughts cannot be cashed out in terms of an ability to recognize that state in ourselves (or in another, for that matter). Lastly, Loar notes that our conceptions of phenomenal qualities sometimes do not involve any sense of independent mastery of a term. This happens, typically, when we think

about some occurrent phenomenal state. We may form various thoughts about it without seeking to apply to it some term that could be independently understood, and against which our application could be deemed successful or unsuccessful. In such cases, it seems odd to call what we have learned an ‘ability’.

The next two chapters will examine, in much greater detail, the phenomenal concept reply to the knowledge argument. (I do not claim to have thoroughly done justice to the ability hypothesis, but space limitations constrain . . .) For now, I wish to turn attention to a related argument, the argument from an explanatory gap.

1.4 The Explanatory Gap

We are directly aware of our experiences. All that we know of them, *qua* experiences, is supplied through attending to them introspectively from a purely subjective point of view. In contrast, our knowledge of their neuroscientific correlates is objective and discursive, rather than subjective and direct. It seems that no matter how advanced our neuroscientific understanding becomes, no matter how perfectly we are able to understand (or simulate) brain function, there will always be an unbridgeable “gap” between such scientific knowledge and the experiential knowledge we acquire by the having of experiences and by attending to them. We will always, it seems, be at a loss to explain why these neural firings are accompanied by the experiences they are, in fact, accompanied by, or why they are accompanied by any experiences at all. Conversely, we seem to be equally at a loss to explain why experiences must have a physical basis.

The argument from an explanatory gap provides a key premise for arguments in favor of irreducible phenomenal properties. Unlike the knowledge or conceivability arguments, the argument from an explanatory gap does not necessarily

purport to *demonstrate* the existence of immaterial qualia; rather it is typically used as a premise for a probabilistic argument to the best explanation (Levine, 1983, 2001). To accept the explanatory gap argument is, at minimum, to hold a) that there is an explanatory gap, and b) that this gap provides evidence for dualism (or, at least, evidence against materialism). (Chalmers, 1996)

There are a number of physicalist responses to the explanatory gap argument. Some deny a gap exists (Dennett, 1991). A few accept the dualist claim that there are irreducibly subjective experiential properties, but hold that some physical properties are irreducibly subjective (Searle, 1992). Some believe there is a gap, but that it need not exist in principle (Nagel, 1974). Others hold that the gap is a product of the limitations of our cognitive architecture, which tells against our ever formulating an adequate theory of phenomenal consciousness (McGinn, 1991).

The proponent of the phenomenal concept strategy, which is the object of our present interest, agrees that the gap is an expected product of our cognitive architecture, but denies that this fact tells against the possibility of an adequate a posteriori physicalist theory of consciousness.

In the section after the next, I will argue that the phenomenal concept strategy has a key advantage over those physicalist theories that take the gap to be closable in principle. The advantage is that in the latter case, but not the former, the gap provides evidence against physicalism. But first, some groundwork.

1.4.1 The Nature and Significance of the Explanatory Gap

There are several models of theoretical reduction on the market. This brief section will lay out the salient features of a simple and influential model. Let us say that a theory T is a set of sentences closed under deduction. (For present purposes, let us assume that the theory is purely extensional). Consider two such theories,

T_1 and T_2 . We may say that T_1 reduces to T_2 just in case:

1) There is a set of biconditional “bridge laws” which connect the distinctive vocabulary of T_1 with the distinctive vocabulary of T_2 (let T_1^* be the theory formed from T_1 when the vocabulary of T_1 is replaced schematically by that of T_2 through application of the bridge laws). 2) All the sentences of T_1^* are sentences in T_2 .

This characterization of reducibility is bare-boned and perhaps too strict. It may be the case that many celebrated accounts of scientific reduction fail to be reductions according to this model. Still, it will suffice for our purposes as a rough schematic of what a theoretical reduction might look like.

Let us suppose that we have some theory M which contains a distinctively mental term such as ‘pain’ and some theory P , which contains a distinctively physical term such as ‘C-fiber stimulation’. Now let us suppose (for the sake of argument) that there is a remarkably robust correlation between feelings of pain and stimulations of neural C-fibers. This correlation allows us to feel quite confident in positing a biconditional bridge law between M and P : namely: *pain* \Leftrightarrow *C – fiberstimulation*. Now it seem that, with some work (and doubtless other bridge laws), we can express all the truths about pain in terms of truths about C-fiber stimulations.

Have we now succeeded in reducing truths about pain to truths about C-fiber stimulation? In a sense, it seems we have. After all, we are supposing not just that the biconditional connecting pain with C-fiber stimulation is true, but that it is a law. And this doesn’t seem *prima facie* unreasonable. If we are convinced that the biconditional is true in a large number of non-actual worlds, then it seems proper to say that we have established a nomic connection between C-fiber stimulation and feelings of pain. As much as anything is a law of nature, we could have it that wherever there’s pain there’s C-fiber stimulation, and vice-versa.

The problem, of course, is that the existence of a nomic connection between C-fiber stimulation and pain is insufficient evidence for the claim that pain is identical to C-fiber stimulation. Identity requires more than nomologically necessary coinstantiation; it requires metaphysically necessary coinstantiation. If pain is C-fiber stimulation then there are no possible worlds, alien as they may be, that contain one and not the other. If there is any possible world in which one is instantiated without the other, then they are not identical.

If pain is not identical to C-fiber stimulation (or to any physical property), then our world is microphysically identical to some other possible world without pain. It is generally taken to be a minimal materialist thesis that all the truths about our world are metaphysically necessitated by the microphysical truths. If the property of being a pain is not identical to any physical property then materialism is false .

To establish the identity of some mental property M and some physical property P we have to show that it is metaphysically impossible for there to be M without P or P without M. It is important to distinguish this metaphysical requirement from the epistemic one. We might not be in the world we think we are in, in which case we could be mistaken in believing that $M = P$. The metaphysical issue ignores this. The question is, if things are as we take them to be, then is it metaphysically impossible for there to be M without P or P without M? Only if this is so are we entitled to claim that $P = M$.

In order to get the strong modal intuition that supports the claim that $M = P$, we need a convincing argument that P necessitates M. This is typically given in the form of an a priori analysis — one that makes it inconceivable, once the relevant facts are known, that there could be P without M.

Let us consider the classic reduction of water to H_2O . We start with a conception of the properties of water. Among other things, it is typically a clear liquid

that freezes when cooled and boils when heated. Now, it is suggested that water is, in fact, just H_2O . What justifies this claim?

Well, one thing that serves to justify the claim is that the water and the H_2O take up the same spatial region. There is a singular concept of spatial dimension that bridges the vocabulary. Secondly, we can see at the molecular level how and why H_2O molecules at room temperature abut one another in a non-rigid pattern. We see that they can slide past one another, and thus understand how a quantity of H_2O is a liquid. Furthermore, we understand that when heated, the mass of H_2O molecules agitates and the most agitated molecules separate and form a gas. We can thus see that agitated H_2O boils, that it *must* boil. We also understand why H_2O , when cooled, forms into a rigid lattice that can only be broken catastrophically. Thus we understand why H_2O must become a solid when cooled.

So what we have here are a collection of concepts (e.g., dimensionality, solidity, boiling, freezing) that bridge the chemical domain and the domain of ordinary perception. Once we make the a posteriori discovery that water is H_2O , we understand how the canonical properties of water are necessitated by the canonical properties of H_2O . Once we make the identification, it becomes impossible to conceive that, for example, the water could be in liquid form while the co-spatial H_2O is in a rigid lattice.

This is not the case with brain states and qualitative states. It seems that no matter how strong a correlation we find between them, we can always conceive that two worlds could be identical in physical respects, yet differ qualitatively. This is because it seems that we lack a set of concepts that can bridge the domains and show this state of affairs to be inconceivable. This inability to generate the required bottom-up necessity is the crux of the explanatory gap.

Although there is widespread agreement that an explanatory gap exists, there

is much disagreement on whether it is bridgeable. Nagel (1974) maintains a certain agnosticism by drawing an analogy between our current state of knowledge of the matter-mind connection with pre-Socratic knowledge of the matter-energy connection. He maintains that we are thoroughly lacking the sorts of concepts that could possibly explain the connection, but does not conclude that we could not, in principle, explain it. Others, such as McGinn, take a more pessimistic view. Human beings are contingent products of evolution with certain sorts of concepts available. It would be hubristic to assume that humans have the conceptual resources to answer all the questions they can pose. Jackson (1982) echoes this pessimistic line.

1.4.2 The Explanatory Gap and Phenomenal Concepts: Reply to Chalmers

A version of the explanatory gap argument formulated explicitly in response to the phenomenal concept strategy has recently been put forward by Chalmers (2006). In this section, I wish to respond to it.

The a posteriori physicalist typically accepts the existence of an explanatory gap between the physical and phenomenal. Her interest in phenomenal concepts therefore lies not in the hope that their study will close the gap, but that it will in some sense explain it or explain away its significance.

There is a kind of explanation, often called “transparent” or “reductive”, which is frequently evoked in the context of the explanatory gap. We have a transparent explanation between a set of low-level truths, (such as physical truths) and high level truths (such as phenomenal truths) if and only if it is inconceivable, given the explanation, that the low-level truths obtain but the high level truths do not.

In the case of the physical explanation of mental phenomena, we can phrase the issue in terms of zombies. Chalmers (2006) writes, “If one can conceive of physical duplicates that lack the key features attributed by thesis C, then there will be an explanatory gap between P and C. That is, there will be no wholly physical explanation that makes transparent why thesis C is true. To explain why, in the actual world, creatures with the relevant physical structure satisfy thesis C, we will need additional explanatory materials, just as we need such principles to explain why actual creatures with this physical structure are conscious.”⁷

Let us suppose that a) a concept is a phenomenal concept only if it refers to a phenomenal property, and b) what makes a property a phenomenal property is not physically explicable. That is, let us suppose that whether or not physicalism is true, there is an explanatory gap between physical properties and phenomenal properties. The proponent of the phenomenal concept strategy will claim that (a) and (b) are both true; a consequence of this is that phenomenal concepts are not transparently explainable in physical terms. To see this, imagine a zombie world, identical to this one in every physical respect, but wherein nothing is conscious. Would the zombies of this world have phenomenal concepts? No.

The reason for this is stipulative (there may be other reasons, but the stipulation is sufficient): for a concept to be a phenomenal concept it must refer to a phenomenal quality. My zombie twin has mental representations narrowly indiscernible from my own, and thus has concepts that are counterparts to my phenomenal concepts. But he has no phenomenal concepts, and thus the concept of a phenomenal concept has not been transparently explained.

Rather than explore some alternative notion of explanation (clearly, there are many), let us move away from the issue of explanation to that of evidence. For it seems that without the right sort of theory of phenomenal concepts, the

⁷(Chalmers, 2006), p. 212

explanatory gap does indeed provide evidence against physicalism. Once such a theory is in place, however, it does not. The reasoning for this is straightforward.

Let us call a physicalist theory that makes no note of the special nature of phenomenal concepts *naive physicalism*. (Naive physicalism is akin to what Chalmers (2001) has dubbed “type-C physicalism”). The view is well-exemplified by Nagel (1974), who writes, “At the present time the status of physicalism is similar to that which the hypothesis that matter is energy would have had if uttered by a pre-Socratic philosopher.”

Now, a naive physicalist takes the explanatory gap to be an unsolved problem, and holds out a non-zero probability that it will be closed — presumably in a way analogous to how previously unexplained gaps, such as between physics and chemistry, have been closed.

So long as we think there’s a chance that naive physicalism is correct, the explanatory gap provides evidence for dualism. Consider the following argument:

- (1) If naive physicalism is true, then it is epistemically possible (for any subject S) that the explanatory gap will not exist at some (arbitrary) future time t.
- (2) Therefore, a rational subject S with non-zero degree of belief in the truth of naive physicalism does not believe with certainty that the explanatory gap will exist at t.
- (3) Dualism entails that the explanatory gap will exist at t.
- (4) Therefore, S’s degree of belief in dualism conditioned on the explanatory gap existing at t is greater than S’s degree of belief in dualism.⁸

⁸Proof: Let $E = \text{prob}(\text{explanatory gap exists at } t)$ and $D = \text{prob}(\text{dualism})$. By definition, $P_E(D) = \frac{P(D \& E)}{P(E)}$. Since D entails E, $P(D \& E) = P(D)$. So we have $P_E(D) = \frac{P(D)}{P(E)}$, and so $P_E(D) > P(D)$, so long as $P(E) < 1$.

- (5) Therefore, at t , the explanatory gap's existing is evidence in favor of dualism for S .⁹

Once we are aware of the relevant features of phenomenal concepts (e.g. their isolation from other kinds of concepts), however, we reject naive physicalism and deny that the explanatory gap might close. The persistence of the explanatory gap is then seen to be evidentially neutral between non-naive physicalism and dualism, both of which entail it.¹⁰

Of course, this argument does not establish that physicalism is true, nor does it bar some arguments to the effect that physicalism is false (such as arguments that rely on a robust link between conceivability and possibility). Nor does this argument defend the claim that there is a physicalistically respectable notion of a phenomenal concept on offer. That is the job for the rest of this thesis.

⁹If we make the further simplifying assumption that S 's degrees of belief for dualism and naive physicalism sum to one, then the explanatory gap's existence at time t is proved to be not just evidence for dualism, but also evidence against physicalism.

¹⁰Proof: $P_E(D) = \frac{P(D)}{P(E)}$ (see footnote 8, above), so if $P(E) = 1$, then $P_E(D) = P(D)$.

Chapter 2

The Demonstrative-Recognitional Account of Phenomenal Concepts

This chapter lays out the type-demonstrative recognitional account of phenomenal concepts and defends it against the charge that we do not possess a large enough supply of such concepts to account for the quantity of phenomenal knowledge we do in fact possess (the next chapter will focus on whether a recognitional account of phenomenal concepts can account for the *quality* of our phenomenal knowledge).

2.1 Type-Demonstrative Phenomenal Concepts

The simplest kind of demonstrative phenomenal concept picks out a *token* experience, that is, a specific instance of a phenomenal property, such as a particular pain occurrence or a particular color experience. By picking out a token, I am able to form thoughts about it, that is, about this particular experience I am having.

Often, experiences are complex. In such cases, I may be able to introspectively point to different *aspects* of my experiences. Sometimes, it makes sense to treat aspects of experiences as experiences themselves. For example, I may have a visual experience that includes both red and green experiences, or an experience that includes both visual and auditory experiences. Other times, it is not so easy to speak of aspects of experiences as themselves experiences (such as the glossiness

of an experience of a color shade, or the particular timbre of an experience of a musical note). It doesn't matter, for present purposes, which terminological choice is made.

By deploying token-demonstratives, I am able to form thoughts about whichever features of experience are salient. I can thus form thoughts of various components of my experiences: I may think that I prefer *this* (the reddish part) to *that* (the greenish part), or *this* (the visual aspect of the experience) to *that* (the auditory aspect of it). The extent to which I can identify and discriminate aspects of experience is a contingent matter that is determined by various capacities, abilities, intentions, and dispositions. The nature and limits of our capacities to deploy token-demonstrative phenomenal concepts determines much of what we can and can't think about and know about our occurrent experiences (Tye, 1995).

Token demonstratives, however, are of limited use to an account of phenomenal concepts. The primary reason is that the fact that I possess a token-demonstrative phenomenal concept of a particular experience (or a particular feature of an experience) does not explain my ability to attend to a particular phenomenal *property*. As token demonstratives refer to singular property-instances, they are of no use in explaining how I am able to recognize various experience tokens as falling under a specific type: how, say, I am able to recognize a particular pain as being qualitatively similar to a pain I've had before, or more generally, to recognize a property instance *as* an instance of the property in question.

The theory of phenomenal concepts that has had by far the greatest impact on recent debates over physicalism is, without a doubt, Brian Loar's (1990/97) "recognitional type-demonstrative" account. According to Loar, a phenomenal concept is a type of primitive mental representation that *demonstrates* a phenomenal property, in the sense that thoughts involving phenomenal concepts typically have the form 'x is one of *that* kind' or 'x is another one of *those*', where x is

an occurrent experience and ‘*that*’ or ‘*those*’ expresses a phenomenal concept that picks out a phenomenal property of which *x* is an instance.

For clarity of exposition, this thesis will not in general distinguish among the several demonstrative-recognitional accounts of phenomenal concepts that exist in the literature, but will rather treat terms such as ‘demonstrative’ and ‘recognitional’ univocally. The exception will be those times when differences are salient to the phenomenal concept strategy or to some antiphysicalist argument.

It should be noted, however, that there are indeed interesting and important differences among the various accounts. For example, demonstrative-recognitional accounts sometimes but not always come with further possession requirements on phenomenal concepts, such as an ability to form images, which are used in the absence of currently perceived instances of a phenomenal property to demonstrate it via deferred ostension (Tye, 2003). This putative aspect of phenomenal concepts is explored in the postscript to chapter three.

There are at least two good reasons for supposing that simple token phenomenal demonstratives are inadequate to the task of referring to phenomenal properties.¹ First, a demonstrated experience will necessarily fall under multiple phenomenal types: for example, a color experience will simultaneously be of some determinate hue and some determinable hue (e.g., a single experience may be both of scarlet and of red). Second, there is the problem of simultaneous occurrences of phenomenal properties.

Type-demonstrative accounts of reference, unlike token-demonstrative accounts (e.g., Tye (1995)), provide a rule for deciding which determinate or determinable is picked out via demonstration. Given this, only concepts that are type- but not token-demonstratives can pick out phenomenal properties. The factor required

¹Perhaps token-demonstrative phenomenal concepts could successfully pick out unrepeatable phenomenal properties. I will not explore this issue.

to underwrite reference to a particular phenomenal property is typically taken to be the subject's ability to re-classify an experience as being of a specific type. Type-demonstrative accounts are thus typically referred to as recognitional accounts, though the term 'type-demonstrative' and the term 'recognitional' differ in meaning.

Loar writes "Phenomenal concepts belong to a wide class of concepts that I will call recognitional concepts. They have the form 'x is one of that kind'; they are type-demonstratives."² However, it is possible to hold that phenomenal concepts are recognitional concepts without to their being type-demonstratives (e.g., McLaughlin/Hill, 1999), and it is possible to hold that there are phenomenal concepts that are type-demonstratives without holding that they are recognitional (e.g., Chalmers, 2003). The issues separate, and it unfortunate that the terms have been often been conflated. In Chapter three, I will give reasons to reject the idea that phenomenal concepts must be demonstratives, though I maintain, with Loar et al, that they are recognitional concepts. An additional consideration against the idea that recognitional phenomenal concepts are demonstratives is that it is plausible that demonstratives always achieve reference via some mode of presentation. This can pose problems for the physicalist in the context of the property dualism argument. I explore this issue at the end of chapter seven.

For physicalists, then, recognitional phenomenal concepts are generally taken to denote those types of brain states (or states of the nervous system) that trigger the tokening of these concepts in certain canonical tasks, such as in an introspective demonstration or re-identification of a phenomenal property.³

²(Loar, 1990/1997), p. 232

³There is a prevailing assumption in the literature that phenomenal properties are tokened locally, i.e., within the subject's neural system, but that does assumption does not seem critical. See (Fisher, 2007), for example, for reasons to doubt the assumption that phenomenal properties are "narrow" in this way.

2.1.1 Thick and Thin Demonstrative Phenomenal Concepts

Chalmers (2006) distinguishes between “thick” and “thin” accounts of phenomenal concepts. Thick accounts, according to Chalmers, propose features of phenomenal concepts that explain our distinctive epistemic situation with respect to consciousness, but which cannot themselves be physically explained. Thin accounts, on the other hand, are those which can be explained physically but do not explain our distinctive epistemic situation with respect to consciousness.

A way to distinguish thick from thin accounts, according to Chalmers, is in the way they handle zombies. A zombie, by stipulation, is a creature that is qualitatively identical to some ordinary human being, but completely lacking in phenomenal consciousness. According to Chalmers, thin accounts will grant my zombie twin phenomenal concepts even as they deny him phenomenal states. Thick accounts, by contrast, will deny my zombie twin phenomenal concepts as well as phenomenal states.

As defined, Chalmers’ taxonomy is quite loaded in favor of certain antiphysicalist arguments, as there is no place in it for phenomenal concepts that are physically explicable *and* explain our distinctive epistemic situation with respect to consciousness. More helpful to us is a dialectically neutral way of putting the distinction: a thick account of phenomenal concepts is “top down” while a thin account is “bottom up”. A thick account will build consciousness into the account from the start, while a thin account will start with purely physical/functional materials and make no assumptions about consciousness.

According to Chalmers, Loar’s account is thick, as it relies on the supposition that a phenomenal state “presents itself” to the subject via non-contingent a mode of presentation, thus enabling “substantive knowledge” of its character. How this

supposition, common in dualist accounts, is to be understood in physical terms is not explained.

While I think that Chalmers' distinction between "top-down" and "bottom-up" theories is helpful, it seems a mistake to jump to the conclusion that my zombie twin would have phenomenal concepts on any thin account. For we may support a theory of phenomenal concepts that "builds in" consciousness only in a *stipulative* way, such as by stipulating it a necessary condition that a phenomenal concept refer to a phenomenal property.

Such an account does not fall easily into Chalmers' thick/thin distinction. It is thick only in the sense that it introduces an unexplained notion of consciousness or phenomenality as a necessary condition on the referent of a phenomenal concept. So according to such an account, a zombie would not possess phenomenal concepts. But it is thin in every other way. It does not "explain our distinctive epistemic situation with respect to consciousness" in the way the dualist would seem to require, if such an explanation requires characterizing phenomenal concepts and their relation to phenomenal properties in ways that go beyond the physical, functional, or causal.

If the demonstrative account of phenomenal concepts is correct, then all that's needed to distinguish phenomenal concepts from other introspectively deployed demonstrative concepts (such as from the introspectively deployed recognitional concepts possessed by those with highly functional blindsight) are differences in denotation. Insofar as we can consistently re-identify a physical property introspectively, our recognitional concept will denote the property in question. And if the referent is a phenomenal property, then the concept in question will be a phenomenal concept. On this account, phenomenal concepts are "thin" in the sense that the referent is determined solely by a physically explicable process, given in terms of dispositional and causal relations between a subject and the

property that is being recognized.

We can think of these type-demonstratives concepts as being representationally primitive in the sense that their tokens do not carry with them information about their referents (unlike, say, discursive concepts). In this sense, they are *characterless*; no information which can serve to identify their referents is intrinsic to phenomenal concept tokens, other than perhaps a simple indexical marker, which serves only to identify various tokens of a phenomenal concept as being tokens of the same concept. Demonstratives achieve reference through a causal chain — to know what phenomenal property a phenomenal concept denotes on a type-demonstrative account of phenomenal concepts, one must trace back the causes that led to the tokening of the concept to isolate the phenomenal property in question.

Loar characterizes his type-demonstrative recognitional phenomenal concepts as “recognitional/demonstrative”, presumably because of the crucial role that recognitional abilities play in their individuation. However, as will become clear, one can (and I argue should) endorse a recognitional account of phenomenal concepts that is not, strictly speaking, demonstrative. I suggest that all the virtues and none of the vices (to be named in chapters four and seven) accrue to a non-demonstrative recognitional account of phenomenal concepts. So I will refer to Loar’s account as “the demonstrative account,” and I will refer to the sort of phenomenal concepts he posits “type-demonstrative” (or just “demonstrative”) concepts.

More needs to be said about these type-demonstrative concepts beyond their classification as demonstratives and their essentially having phenomenal properties as their referents. For one thing, they need to be understood as demonstrating phenomenal properties in a particular way — via an “introspective perspective,”

and a physicalists cannot simply present this perspective as a bare relation between a subject and a phenomenal property. Rather, the relation will have to be explicated in a physicalistically acceptable way.

A key requirement on this introspective relation is that such a relation, of necessity, can obtain only between a subject and a phenomenal property she is having. Also, the causal connection cannot be deviant. The subject must recognize the phenomenal property in virtue of her having it, and not in virtue of any inessential effects of the having of it. They must be perspectival, not discursive.

Such a requirement is necessary to forestall numerous counter-examples to a recognitional account of phenomenal concepts. Imagine, for example, a subject whose pulse quickens whenever she has a visual experience that includes the color red, and that this particular kind of pulse quickening (say, the acceleration of the pulse rate and its final level) occurs only when she has such an experience. We can imagine that she is able to reliably detect the color experience by reliably detecting the particular type of change in pulse it causes. We can further suppose that she is entirely unaware (at the cognitive level) that she is recognizing changes in her pulse rate. So her demonstratives are “bare”, in the sense that they are not mediated at the cognitive level by any discursive mode of presentation. Yet it seems that they would not count as phenomenal concepts, as they seem to aim at the wrong target. They are not, it seems, being introspectively applied in the right way.

The demonstrative account will have to give an account of introspection that explains why certain links in the causal chain between the tokening of phenomenal property and the tokening of a phenomenal concept count are acceptable but others are not. After all, there must be some intermediate links in the causal chain between the tokening of a phenomenal property and the tokening of a

phenomenal concept of the phenomenal property.

Of course, those who reject the demonstrative model may well reject the supposition that there are causal intermediaries that stand between a tokening of a phenomenal property and a tokening of a phenomenal concept of it. Some, predominantly dualists, agree with the demonstrative theorists that the phenomenal property and phenomenal concept tokens are distinct, but hold that they are connected by a relation more “intimate” than causation, such as “acquaintance”. Others claim that phenomenal concepts contain or “quote” their referents, and for this reason there is no room for causal intermediaries. These rival views are considered in chapter three of this thesis.

2.2 Do We Possess Enough Introspectively Applied Type-demonstrative Concepts to Account for the Extent of Our Phenomenal Knowledge?

It is not a problem for the demonstrative theory of phenomenal concepts that our stock of demonstrative phenomenal concepts is vastly outstripped by the stock of phenomenal properties we’ve experienced. For it is reasonable to suppose that we cannot, or at least do not, conceptualize all the phenomenal properties (and aspects thereof) we have.

This seems commonsensical. To take it to be problematic is, in essence, to assert that abilities or capacities to conceptualize phenomenal properties figure criterially into the having of such properties. But such a view is unmotivated.

While there may well be contingent facts that connect the ability to have a phenomenal property with the ability to conceptualize it, to assert that such dependencies are necessary is close to asserting that phenomenal facts are to be cashed out in terms of conceptual facts. Such would be expected of a kind of

reductive account of consciousness, according to which the having of a phenomenal property just is, in whole or in part, the having of certain conceptualizations or the having of certain abilities to conceptualize. While such theories exist, they are not popular nowadays.

It is beyond the scope of this work to argue against reductive theories of consciousness (what Chalmers (2001) calls ‘Type-A’ theories of consciousness). Rather, the present debate over the nature of phenomenal concepts should be understood as being a debate between those physicalists who do not think that truths about phenomenal consciousness are reducible a priori to non-phenomenal truths (Chalmers’ Type-B physicalists), and those critics who assert that such an a posteriori physicalism is an inconsistent position.

That being said, it is a problem for the demonstrative theory if (a) “knowing what it’s like to have Q” is to be cashed out in terms of possession of a type-demonstrative phenomenal concept of Q and (b) our stock of type-demonstrative phenomenal concepts is outstripped by the stock of phenomenal properties that we *do* know what it’s like to have.

In a prominent critique of type-demonstrative accounts of phenomenal concepts, Diana Raffman (1995) argues, using empirical studies as evidence, that we do not possess nearly enough type-demonstrative recognitional phenomenal concepts to account for our phenomenal knowledge. If true, this spells trouble for the type-demonstrative theory of phenomenal concepts.⁴ Raffman’s objection is similar to one often raised against the ability hypothesis in response to inverted spectrum worries — the fear is that we do not possess a large enough stock of phenomenal concepts to capture introspectable differences among experiences.

Raffman points to several studies that suggest that normal subjects can detect

⁴Raffman holds, in addition, that her argument spells trouble for physicalism more generally. We will not explore that claim here.

just noticeable differences between more color shades in simultaneous presentation than they are capable of re-identifying over time. This difference, furthermore, is not small. In fact, it seems that the ordinary subject's ability to discern shades in simultaneous presentation vastly outstrips her ability to re-identify shades. If this is so, then the physicalist is presented with a dilemma, or so Raffman argues.

If it is granted that the subject knows what it's like to see each of the discriminable shades presented, then it follows that knowing what it's like cannot consist in possessing a type-demonstrative concept. So, Raffman reasons, the first horn is to grant that the possession of token demonstratives must be sufficient for accounting for what it's like to have a phenomenal property. Although we may possess discursive concepts of these phenomenal properties that we can discriminate but not re-identify, such concepts are not "direct" in the sense that the physicalist would like and the dualist demands (i.e., the reference must not change across possible worlds). For the reasons we've discussed earlier, token demonstrative concepts are not up to the task of distinguishing among phenomenal determinates and determinables. The second horn is to reject physicalism.

If acquiring phenomenal knowledge is cashed out in terms of acquiring a type-demonstrative phenomenal concept, then there are multitudes of experiences that we can have, attend to, but not know what it's like to have. But, Raffman argues, this is counter-intuitive. If Jackson's Mary were presented, say, with an experience of red-31 (a shade off unique red that cannot be re-identified at the specificity level of a just noticeable difference), she'd assert, while attending to the shade, that she now knows what it's like to see red-31. It would seem quite counterintuitive to suggest that Mary is wrong here.

The point becomes more acute when we contrast this situation with one in which the color that Mary is presented is not red-31, but unique red, a shade which subjects typically can re-identify and distinguish from just noticeably different

shades in non-simultaneous presentation. In this latter case, the supporter of the type-demonstrative account of phenomenal concepts would now hold that Mary is correct when she exclaims (while attending to the shade) that she now knows what it's like to see unique red.

But surely it seems (so the argument goes) that the two situations are identical with respect to Mary's having or lacking phenomenal knowledge of the specific shade she's currently seeing. After all, how could her ability to re-identify unique red but not red-31 have any bearing on her knowledge (or lack of knowledge) with respect to the shade she is currently attending to?

I think the first thing to be said is that this is setting the burden of proof on the wrong side. There is no a priori reason to suppose that there will be symmetries such as the antiphysicist is supposing. Cognitive psychology has shown there to be numerous (if not always salient) differences in how we attend to and discriminate various colors.

The physicalist, it seems to me, should deny this supposedly obvious intuition that Mary's phenomenal knowledge is the same in the case of red-31 and unique red, even as she is having the shade experience in question. There are several ways to deny this. One is to claim that there are multiple senses of "knowing what it's like" in play. In one sense, possessing an appropriate token demonstrative is sufficient. Anyone who can discriminate shades in simultaneous presentation and form thoughts about the objects of such discrimination possesses some phenomenal knowledge, which one who cannot perform such discrimination lacks. And perhaps it is this sense of "knowing what it's like" that is salient in the case of Mary's initial reaction upon being presented with a color sample for the first time.

There are, however, other questions Mary might ask of her newfound phenomenal knowledge. She might, for instance, wonder whether she is, in fact, conceptualizing a particular shade or some broader swath of the color spectrum.

And this might be puzzling to Mary, as she will not know the answer to this question so long as she is having her initial color experience.

It is thus reasonable to suppose that Mary's (alleged) initial assertion that she knows what it's like to experience the particular shade she is experiencing might turn more circumspect. Rather than continue to believe that she knows what it's like to see red-31, she may well come to believe, rather, that she knows what it's like to see some determinable of the particular shade she picks out with her introspectively applied phenomenal token-demonstratives.

Mary may even come to the belief that there is no way for her to know whether she knows what it is like to see this particular shade or some determinable of it, and she may devise tests for herself to determine whether she is conceptualizing a particular shade or some swath of the spectrum. And these tests may well involve tests of her abilities, including abilities to re-identify phenomenal properties in non-simultaneous presentation.

This is, of course, just what the proponent of recognitional phenomenal concepts demands, and such a proponent should insist that it is by no means obvious that Mary would not demand such tests (or the knowledge of what the results would be of such tests) before confidently asserting that she knows exactly which phenomenal property she has gained phenomenal knowledge of. The bare token-demonstrative might well be insufficient to pinpoint the relevant feature (as token demonstratives, in general, are).

The upshot is that one can acknowledge the cognitive similarities between the cases of Mary's having an experience of unique red and her having an experience of red-31, as well as Mary's intuitions in both cases that there is a sense in which she knows what it's like to see the shade she is seeing. But the physicalist should cash out the similarity in terms of Mary's salient concerns, rather than in blanket terms regarding her phenomenal knowledge.

While the color sample is present, the focus is on token-demonstration and the ability to distinguish the shade from others in simultaneous presentation. But this concern will pass as Mary moves to consider the more substantive sense of knowing what it's like that demands the ability to re-identify the phenomenal property in question. When her salient interests change, she may well lose the intuition that she is sure that she knows (or knew) what it is like to see red-31.

One reason to suppose that Mary's opinion might change once the sample is removed is that we often do not know the scope of our abilities until they are tested. For example, an actor may not know whether he's memorized his lines until he tries to recite them without a script. Similarly, Mary may well not know something crucial about her knowledge (or lack of knowledge) of red-31 until her experience of it is concluded.

To be clear, the claim is not that there is one sense of knowing what it's like that applies to occurrent experiences and another sense that applies to phenomenal properties one is not currently having. Mary could, of course, be interested in the more demanding sense of knowing what it's like, i.e., the kind that requires possession of a type-demonstrative phenomenal concept, from the beginning. Rather, the suggestion is that the intuition which motivates the claim that Mary knows what it's like to have an experience of red-31 illicitly conflates the relatively uncontroversial claim that Mary, upon release, acquires a variety of concepts of red-31. Some are discursive and some are demonstrative. But I suggest none are sufficient for knowing what it's like. And it does not seem obvious that Mary, upon reflection, would not agree.

Chapter 3

Acquaintance, Revelation, and the Quotational Account of Phenomenal Concepts

3.1 Chapter Introduction

An influential and diverse group of philosophers has expressed dissatisfaction with demonstrative-recognitional accounts of phenomenal concepts, holding that such concepts cannot account for the *nature* or *quality* of the phenomenal knowledge that is often gained from the having of an experience. While the specific complaints vary, there are common threads. Although it is granted by these objectors that introspectively applied recognitional abilities may play a role in phenomenal knowledge, it is their contention that there must be more to the story. The worry sometimes manifests itself as a product of phenomenological reflection: it just seems that there is more to deploying a phenomenal concept than baldly demonstrating a phenomenal concept. It seems, for example, that when thinking about the nature of experiences of red, I am having a *substantively* different thought (not just referentially different) than I have when I think about experiences of green.

Beyond phenomenological reflection, worries arise in response to the existence of *non*-phenomenal properties that we can introspectively recognize. There are introspectively applied demonstratives that do not denote phenomenal properties.

Cases of “blindsight”, for instance, are often brought up as putative counterexamples to the recognitional account of phenomenal concepts.¹ Thus, the objection continues, to characterize phenomenal concepts as introspectively acquired demonstratives doesn’t capture what is special about them or what distinguishes phenomenal knowledge from ordinary recognitional knowledge. And so, if phenomenal knowledge is to be cashed out in terms of the possession of a certain kind of concept, that concept cannot be a demonstrative concept (or any other sort of merely recognitional concept).

These sorts of concerns have prompted some noted philosophers, including some physicalists, to posit additional constraints on phenomenal concepts. One suggestion is to revive and refurbish the classic notion of *acquaintance*, which is a relation that is taken to be more “intimate” than a mere causal relation. Dualists typically take acquaintance to be a basic relation between a subject and a non-physical phenomenal property (or a token of a phenomenal property), and physicalists attempt to spell out the acquaintance relation in physicalistically respectable terms. For example, Chalmers (2003), a dualist, Levine (2006), an agnostic, and Balog (2006), a physicalist, all seem to agree that understanding the nature of phenomenal knowledge requires the exposition of some sort of special relation of acquaintance. A related suggestion is to hold that phenomenal concepts are in some sense partially constituted by, or “quote”, samples of their referents. Such accounts have been proffered by, among others, Chalmers (2003), Balog (2006), Papineau (2002) and Block (2006).

In this chapter I will argue that the recognitional account of phenomenal concepts should not be abandoned or substantially modified, at least not in response to the sorts of worries that lead theorists to posit a special acquaintance relation

¹Although blindsight gets much of the attention in the literature, it seems to me that less exotic cases abound, such as introspective awareness of non-phenomenal proprioceptive properties (e.g., my awareness of my elbow’s being bent).

or a quotational account of phenomenal concepts.

To explore what has seemed to some to be missing from recognitional accounts of phenomenal concepts, I will discuss the phenomenon of blindsight, and what sorts of perceptual and/or experiential concepts blindsighted subjects possess. I will then consider more phenomenologically motivated arguments against the recognitional account. I will examine the intuition behind the desire for a basic acquaintance relation and attempt to explain it away. Finally, I will argue that the idea that phenomenal concepts contain or quote their referents can be understood in two very different ways.

I suggest that on one reading of the quotational account it is highly implausible that any of us in fact possess concepts of the kind in question, and that on another reading the quotational account is functionally indistinguishable from a more orthodox recognitional account. If I am right, then quotational accounts are either highly implausible or functionally indistinguishable from recognitional accounts. In neither case, then, does a quotational account of phenomenal concepts present an advantage over a recognitional account.

3.2 The Challenge of Blindsight

Blindsight is a phenomenon in which people who are perceptually blind in a certain area of their visual field demonstrate some response to visual stimuli, without any associated qualitative experience. In Type 1 blindsight, subjects have no awareness whatsoever of any stimuli yet are able to predict, at levels significantly above chance, aspects of a visual stimulus, such as location, or type of movement, often in a forced-response or guessing situation. Type 2 blindsight is when subjects have some awareness of, for example, movement within the blind area, but no visual percept. This may be caused by, for example, the subject

being aware of her eyes' tracking motion, which functions normally. Both types of blindsight have been thought to present problems for recognitional accounts of phenomenal concepts, but I will confine the discussion to Type 2, as it presents a clearer challenge to the accounts.

The salient purported difficulty is that blindsighted subjects, particularly highly functioning ones (notably those of Type 2), possess abilities to re-identify visual perceptual occurrences, even in the absence of associated qualitative experience. They thus meet the necessary functional requirements for the possession of recognitional concepts, at least insofar as ordinary sighted individuals meet the necessary functional requirements for the possession of recognitional phenomenal concepts. They meet these requirements not simply because they are able to re-identify particular perceptual phenomena, but because, arguably, they are able to do so "introspectively", in a way very similar to that which has been presumed to serve to distinguish recognitional phenomenal concepts from other sorts of recognitional concepts.²

The concepts deployed by the blindsighted are recognitional, introspectively applied concepts, and yet are not phenomenal concepts. So, the worry goes, phenomenal concepts cannot be distinguished from other sorts of concepts merely by their introspectively-deployed recognitional nature. And if this is the case, then it seems that what essentially makes a phenomenal concept a *phenomenal* concept has eluded the recognitional account of phenomenal concepts. And so, the analysis concludes, some different or further substantive requirements on the nature of phenomenal concepts must obtain. Or, alternatively, a special relation that obtains between subjects and the phenomenal qualities they token, but not

²For example, the fact that recognitional phenomenal concepts are introspectively applied distinguishes them from, say, the sort of recognitional concept I might hypothetically acquire from being able to read a brain-scan, and recognizing my own pain occurrences via say, looking at a self-applied scanner and noticing C-fiber stimulation occurring.

between the blindsighted and the contents of their perceptual concepts, must be spelled out as a part of an adequate account of phenomenal knowledge.

To this sort of criticism, I suggest the proponent of a recognitional account of phenomenal concepts take a hard-line: *What serves to distinguish phenomenal concepts from those recognitional concepts deployed in certain blindsight cases is the nature of the referent and nothing more.* That is, it is sufficient to note that it is a requirement on phenomenal concepts that they denote phenomenal properties. This requirement is met in the case of ordinary visual perception (which involves a perceptual experience) but not in the blindsight case.

A response similar to mine has been recently articulated by Janet Levin.³ She writes:

In particular, if phenomenal concepts really function like introspectively deployed demonstratives, then all that's needed to distinguish them from introspectively deployed nonphenomenal demonstratives are differences in what they denote. This is the standard way to think about demonstratives deployed from similar perspectives: if my (token) "that" (pointing out the window) picks out a car with a particular constellation of properties, then what determines whether I've denoted a Maserati (and not the nearby Ford), and if my "that (kind of car)" consistently picks out all and only Maseratis, then I've managed to denote the (very distinctive) property of being a Maserati (rather than being a Ford). My demonstratives may thus be regarded as special (after all, they're Maserati demonstratives), but this is entirely because of the features of what they denote. Similarly (focusing here on token demonstratives), the states a blindsight subject's demonstratives denote are not experiences, since these states (by hypothesis) have no phenomenal properties, whereas the states denoted by a normal person's demonstratives do have phenomenal properties. If phenomenal properties are identical with physical properties, then there will be physical differences between our experiences and the introspectively denoted neural states of a blindsight subject. Some of these physical properties, of course, will be "felt" by the subject who have them, and some will not — but that's just

³Levin's paper was published in late 2006, after I had formulated the bulk of this chapter in the summer of that year. I had not seen a manuscript of her paper prior to publication, which is a shame, as she has eloquently articulated a position similar to my own — at least in regard to how we respond to demands for an acquaintance relation and to suggestions that phenomenal concepts are quotational. Our positions diverge considerably, however, with respect to whether phenomenal concepts are demonstratives. She (tentatively) holds with Loar that they are. I hold that phenomenal concepts are non-demonstrative recognitional concepts (See chapters four and seven of this thesis).

what it is for some, but not all, of them to be phenomenal. And that's all it should take for the demonstratives that denote these states, or their associated phenomenal properties, to be "special" as well.⁴

This "hard-line" response may not be quite so hard as it first appears. It does not, for instance, deny that there are functional differences between phenomenal concepts and non-phenomenal introspectively deployed recognitional concepts. For example, subjects will tend to classify the referents of phenomenal concepts as determinates of a variety of determinables, such as "color experiences", "visual experiences", "perceptual experiences", or simply "experiences". Blindsighted subjects too will know that the denotations of their auditory, olfactory, and tactile phenomenal concepts have something in common, i.e., phenomenality, which the denotations of their non-phenomenal visual perceptual concepts lack.

In other words, subjects are typically able to recognize experiences as experiences, and thus are able to functionally demarcate phenomenal concepts from non-phenomenal introspectively deployed recognitional concepts. Thus there is some functional substance (that is, a cognitively discernible difference) behind the seemingly stipulative divide between phenomenal and non-phenomenal introspectively deployed recognitional concepts.

3.3 Quotational Accounts of Phenomenal Concepts

On the dualist side, Chalmers (2003) posits what he calls "pure" phenomenal concepts (in addition to the type-demonstrative phenomenal concepts mentioned earlier), in which the referent is "present" in the sense of the concept. Some physicalists have similarly suggested that phenomenal concept tokens contain tokens of the properties they refer to (Papineau, 2002, Balog, 2006, Block, 2006). According to Papineau, "Phenomenal concepts are compound terms, formed by

⁴(Levin, 2006), p. 94

entering some state of perceptual classification . . . into the frame provided by a general experience operator ‘the experience: —’. Such terms will have a sort of self-referential structure. Very roughly speaking, we refer to a certain experience by producing an example of it.”⁵

One of the first things to note about the quotational aspect of the account is that it is insufficient to determine the referent of a phenomenal concept. If a phenomenal concept of Q were to refer to Q just in virtue of producing a token of Q, there would be no way to distinguish between a phenomenal concept of Q and a phenomenal concept of a determinable of Q. Plausibly, both concepts could be formed by entering a token of the determinable into the operator ‘the experience: —’. So, according to the quotational account, both concepts would be identical: they would refer to concepts that resemble that token. But this is of course false: we have two different concepts, one refers to Q and the other refers to the determinable of Q. (This worry is, of course, similar to that raised for a token-demonstrative account of phenomenal concepts).

Papineau realizes that the quotational character of phenomenal concepts, on his view, is not sufficient for fixing the reference of such concepts. He argues that we would need to supplement such an account with an additional semantic theory which explains how a given phenomenal concept can refer to the corresponding phenomenal property (a token of which is incorporated into the concept).

Papineau suggests that we can appeal to theories of content such as Fodor’s causal account or Millikan’s teleological account. Phenomenal concepts, according to Papineau, refer to items that resemble their ‘fillings’ because applications of these concepts are typically caused by those items, or because it is the function of such concepts to track those items.

⁵(Papineau, 2002), p. 116

According to the recognitional account, what fixes the reference of a phenomenal concept is our disposition to recognize tokens of a certain property. This too is a causal theory: the concept refers to that property, tokens of which causally trigger the concept to be applied (in appropriate circumstances). So there does not appear to be anything to recommend one theory over the other as far as reference is concerned.

I suggest that the quotational account is either implausible or vacuous: that depending how the account is spelled out, either we do not possess quotational phenomenal concepts or such concepts have no explanatory value. But first, to set the stage, let us consider how the account addresses an obvious issue. Generally speaking, when we think about, remember, or imagine some phenomenal property Q, we do not detect Q in the experience of thinking, remembering, or imagining. When we think imaginatively about what it's like to feel some type of pain, for example, we do not thereby experience a pain of that type. This is not to say we don't experience some discomfort at the thought. The point here is that we don't feel the pain we were thinking about. So, assuming we are not simply failing to notice experiences of Q, it follows that a phenomenal concept of Q can be deployed in the absence of an experience of Q. Granted, it seems that we are having a distinctive sort of experience when thinking imaginatively about some pain, one that differs from what we experience when thinking imaginatively about, say, some tickle. The point here is a modest one: *we don't generally have the type of experience we're thinking about.*

A proponent of the quotational account may say that although we may not have an experience of the phenomenal property *itself* in such cases, we have an experience of an *image* of it, which, while different, resembles it in some appropriate way. There are two questions to ask of this response. First, do we *always* token an image when thinking about what it's like to have some type of

experience (except when actually having an experience of that type)? Second, if so, does this fact motivate the quotational account?

I take the answer to the first question to depend on a conceptual analysis of the expression ‘thinking about what it’s like’. On the face of it, it seems that I needn’t have an experience of Q nor an experience of the image of Q when I have a thought about what it’s like to experience Q, for example when I truthfully think “I am not now experiencing nor imaginatively recreating a sensation of pain.” The proponent of the quotational account must hold that my thought, if true, does not involve a phenomenal concept of pain or does not involve the kind of phenomenal concept that’s relevant to knowing what it’s like. (If the proponent of the quotational account takes this latter position, he should provide reason to suppose that the quotational account has any advantage over a recognitional account, supplemented with the requirement that Q or an image of Q be tokened concurrently with a phenomenal concept of Q. I see nothing to recommend restricting phenomenal concepts this way, but if someone wishes to insist on it, I suggest supplementing the recognitional account in this simple fashion).

Even if we were to accept that an image must be experienced when thinking thoughts relevant to knowing what it’s like, the quotational account would face serious problems. Imaginative re-creations of phenomenal properties often provide far less detail than can be discerned and thought about when the properties themselves are instantiated in experience. We are, for example, notoriously bad at imaginatively re-creating experiences of odors, even as we are exceptionally good at introspectively discerning and characterizing fine distinctions among them (Köster, 2005).

On top of this, we often imagine phenomenal properties *wrongly*, even those we take to be simple and well-known. To take an example, think about what it’s

like to see black. In doing so, you may well have imagined a uniform color field (such as what you'd experience in complete darkness). The problem is that what you imagined was not black, but brain gray. Black is a contrast color, and cannot be experienced on its own. But I take it to be uncontroversial that we know what it's like to have an experience of black.

That imagination provides neither the detail nor accuracy of introspective awareness is not problematic in itself (one needn't hold that knowing what it's like to have an experience entails being able to form a good image of it), but it suggests a dilemma for someone who insists that thoughts involved in knowing what it's like require the experience or an image of it: *either the levels of detail and veracity of the image don't much matter (in which case, why believe it to be a necessary component of such thoughts?), or we can't really think about what an experience is like except when having it.*

One could perhaps retain a use for the quotational account by holding that knowing what it's like to have an experience of Q consists in having an ability or disposition to think certain true thoughts about Q while having an experience of Q, where such thoughts involve quotational phenomenal concepts. Such a view would grant quotational concepts a critical epistemological role without saddling them with the problems we've raised. The position is interesting, but one would need to accept, implausibly in my view, that one could know that p without being able to think that p (except in highly specific experiential circumstances).

There is a further problem for the physicalist who takes this position. One might wonder what purpose is served by including the phenomenal property in the concept. An account will still have to be given as to how, when the concept is tokened, the included property instance interacts with the rest of the tokened concept or thought. It is doubtful the physicalist can explain this interaction in a way that is more plausible than causal interaction. Given this, it would seem

simpler to eliminate quotation and stipulate that knowing what it's like involves the ability to token a type-demonstrative phenomenal concept while its referent is instantiated. Quotation, in the end, presents some implausibilities that, when adequately addressed, renders the quotational aspect of the account explanatorily moot.

3.4 Awareness and Acquaintance

Our knowledge of our experiences seems so immediate and complete that one might find odd the idea that we can explain phenomenal knowledge without positing a basic, intimate relation such as acquaintance, through which the *essence* or *intrinsic nature* of phenomenal properties is revealed. In this section, I wish to dispel this common misgiving.

3.4.1 An Innocent Notion of Awareness

We are often *aware* of phenomenal properties we instantiate in experience. This awareness is *direct*, in the sense that it depends neither on awareness of anything else nor on any process of inference. We are at first pass aware *of* instantiations of phenomenal properties, not *that* anything is true of them. Awareness of an experience thus does not by itself provide knowledge of what it's like to have it. (I assume the prevailing view that phenomenal knowledge is comprised, at least in part, by knowledge of facts (Tye, 2000), (Chalmers, 2003), (Loar, 1990/1997).

While awareness does not by itself provide factual knowledge, it is nevertheless the starting point for much of it. We are typically able to learn truths about those aspects of experience of which we are aware, provided we are able to attend to them and form appropriate concepts to stand for them in thought. A subject who acquires a phenomenal concept of Q on the basis of awareness of Q typically

thereby *knows what it's like* to have Q.

It would be nice to use the term 'acquaintance' to describe this relation. However, I will use the term 'awareness' in order to avoid any unhelpful ambiguity in this discussion. It might, however, be more elucidating to see the present project as presenting and defending a more modest conception of acquaintance, rather than as a rejection of acquaintance outright.

Awareness stands in innocent contrast to acquaintance, at least as the term 'acquaintance' has been used since Russell. According to Russell, we have knowledge by acquaintance of anything with which we are directly (i.e., non-inferentially) aware. Of interest to us are (token) conscious perceptions. We may also add memories, thoughts, emotional states and imaginings to the list. According to Russell, we gain acquaintance-based knowledge through deliberative introspection. Not only am I acquainted with, say, my toothache, but I am (typically) acquainted with the acquaintance. As long as they are objects of conscious experience, they are known by acquaintance. This is clearest in the case of perceptual (and other experiential) universals. We know redness, or pain, not from any chain of inference, but by intuiting redness and pain from red patches and painful episodes. We come to be acquainted with emotions and all other experiential universals in a similar fashion.

Now Russell, and others, have had particular views of the structure and content of sense data and of the epistemological castles that could be built on their foundations, but that is irrelevant to our present purposes. The important thing to note here is that acquaintance events are what Russell called sensations. Russell is careful to draw a sharp distinction between the acquaintance act, and that with which the act acquaints us.

Consider an experience of pain. According to Russell, this experience gives me knowledge by acquaintance of pain. There is no distinction drawn between

the experience and the acquaintance act. The acquaintance consists just in the experience — there is no acquaintance without experience, and there is no experience without acquaintance. Furthermore, in no way does Russell present this as a contingent fact. The acquaintance act and the experience are semantically tied in the following way: acquaintance is just the epistemic character of experience, and all experience has epistemic character.

Could an identity theory work in Russell's framework? I don't think so. Suppose we wanted to identify pain events with, say, C-fiber stimulation events. According to Russell, pain events are acquaintance events — they are experiences of something. By substitution of identicals, we would have it that C-fiber stimulation events are experiences of something. But of what? The ready answer seems to be *pain*, but now we're just chasing our own tail, so to speak.

There is, of course, an ambiguity in a phrase such as 'experience of pain'. It can be seen as asserting either the existence of one object (i.e., an experience) or of two (i.e., an experience and a thing experienced). Russellian acquaintance requires the latter of the disambiguations, but this view is not conducive to physicalism.

We can see from Russell's account that the driving force is epistemological. Russell's thesis that an experience just is an acquaintance act (i.e., a kind of knowledge act), and vice versa, has few adherents today. It is common coin among theorists to suppose that an experience of Q can occur without a subject's thereby gaining knowledge of Q. Something other than (mere) epistemology is driving the renewed interest in acquaintance.

With the renewed interest in acquaintance (that is, an acquaintance stripped of Russell's epistemological demands and ties to sense-data), one might wonder why the fuss? What about it appeals, and what ties it to quotational accounts of phenomenal concepts?

Chalmers writes:

In the prima facie justification of a direct phenomenal belief by an experience, there are three factors of the same sort. First, content: the epistemic content of the direct phenomenal belief must mirror the quality of the experience. Second, a natural connection: the phenomenal belief must be appropriately constituted by the experience. And third, epistemic status: the subject must be acquainted with the justifying quality. The details of the requirements are different, as befits the difference between belief and experience, but the basic pattern is very similar.⁶

It is the second of these requirements that will be of interest here. What ties together the demand for acquaintance with the suggestion that phenomenal concepts quote their referents, is the idea that phenomenal concepts (and phenomenal knowledge) are in some sense ontologically dependent on their referents. To think about a phenomenal property (or property instance, in Chalmers's view), is to in some sense be privy to the very nature of the thing under consideration. We might put this in terms of a thesis of *Revelation*, which can serve to motivate both the epistemological doctrine of acquaintance and the ontological doctrine of quotation.

3.4.2 The Revelation Intuition

Revelation. *A subject who acquires a phenomenal concept of phenomenal property Q on the basis of awareness of an instance of Q typically thereby knows the essence of Q .*

There is prima facie evidence against Revelation. Hue qualities, for example, may appear to be simple but have essential phenomenal aspects which may not

⁶(Chalmers, 2003), p. 241

be noticed. Binary hues, for instance, are phenomenally complex: e.g., it is essential to an experience of orange that it appears both reddish and yellowish. But surely someone needn't know this to know what it's like to see orange. Even unitary color experiences have essential phenomenal aspects (i.e., aspects of hue, saturation, lightness) that may escape notice (To be clear, we're talking about phenomenal aspects of color experience, not properties of colors themselves).

The same goes for other modalities. Experiencing various overtones and undertones is essential to the experience of hearing a note from a piano, however simple the experience might seem. The upshot is that an experience may be phenomenally complex, even as we represent it in thought via a primitive phenomenal concept (more will be said on this in the last section of chapter five). Nevertheless, we know what it's like to have these experiences. The ignorance goes deeper when we consider that our abilities even to recognize various types of experiences often comes gradually, through repeated exposure or deliberate effort.

Frequently, we can enhance our recognitional abilities, both through voluntary and involuntary means. Voluntarily, we can assimilate and apply theoretical information about where a certain shade lies in quality space. On this approach, the re-identification is not immediate, but rather occurs via reference to a target shade, such as unique red, which is readily re-identifiable to most subjects (see chapter two).

This technique will be familiar to those musicians who, having been accustomed to a tuning note (typically an A), develop the ability to recognize the note. By applying various techniques, one can then move from that note via a set interval to the note of one's choice. The capacity for successful employment of such techniques will, of course, vary by individual.

Should possession of such abilities count in favor of possession of the associated recognitional phenomenal concepts? It seems not. Although the subject does

possess certain recognitional abilities, it seems her relation to her experience is not “direct” enough to count. Put more explicitly, in such cases reference is not achieved via descriptively unmediated demonstration, but rather via a hybrid process involving demonstration and descriptive mediation. Were this the end of the story, no inroads would be made against intuitions supporting Revelation.

However, we can readily imagine someone coming to be able to recognize new types of experiences consistently, without recourse to the descriptively mediated process. In such cases, the re-identifications one learns to make no longer depend on association or inference, even though such associations and inferences were causally responsible for the acquisition of the newly acquired recognitional abilities. As Levin (2006) notes, there are many familiar cases — distinctive kinds of vegetation, fabrics, styles of painting, or even unusual shapes — in which it takes time and instruction to acquire the ability to reidentify various phenomena, and thus the type-demonstrative concepts of those phenomena.

In such cases, one might use association or inference to draw conclusions about the way things are from the way they seem: fabrics that feel like this are likely to be silk; paintings that look that way are likely to be Cezannes. But coming to acquire the recognitional concepts themselves doesn’t ultimately depend on the learning of various inferences, but rather on learning to attend to those distinctive features of the way things seem that qualify them as seeming that way.

3.4.3 There are Numerous Ways to Acquire Phenomenal Concepts

Though those phenomenal properties we can introspectively discern in simultaneous presentation will always outnumber the properties for which we can have

recognitional phenomenal concepts, it's also a psychological fact that we can increase the range of these recognitional concepts, at least to some degree through various sorts of instruction and practice.

For example, people who have taken a course in music appreciation or wine tasting are able to recognize, or identify consistently, sounds or tastes that they were never able to recognize before, despite being able to discriminate them in simultaneous presentation.

It seems reasonable to label what subjects acquire from these courses newfound phenomenal knowledge of the various types of experiences associated with music or wine.⁷ What someone gains from experience, in coming to know what it's like to have it, is — as the type-demonstrative theorist suggests — a new set of concepts. These concepts are distinct from any that she already possesses. They denote directly, as demonstratives. There is no need to invoke murky and unexplanatory notions, such as a tighter-than-causal acquaintance relation.

Once the prima facie plausibility of Revelation is dispelled, the purported need for acquaintance to explain the “substantive” knowledge experience provides vanishes. *In at least some (and maybe most) cases, one has to know a lot about the properties in question, in particular their interrelations with others in various quality spaces, to acquire the recognitional concepts — a phenomenon that accounts for the intuition that coming to know “what it's like” is coming to know something substantive.*

Still, one may object that this view is not plausible in the least. What happens when we acquire new recognitional concepts is something much less cognitively

⁷Levin (2006) notes that if the ability to discriminate among items (or experiences) with different properties, when presented simultaneously, itself counts as having substantive knowledge, or cognitive grasp, of those properties, it's hard to describe what people in courses such as these have learned. For this reason, too, the intuition that merely having an experience, and presumably being acquainted with it, provides substantive knowledge of the property experienced can be dispelled.

complex, something that doesn't require anything like these explicit comparisons or other applications of theory that have been described. What goes on in wine-tasting or music-appreciation classes — let alone what goes on when our recognitional dispositions are “naturally” shaped and enhanced by our interactions with the world, simply does not resemble the sort of deliberative learning via comparison that goes on, for example, in wine tasting courses, when one is shown an example of the target item (and perhaps a foil that's quite different along a certain dimension) and is told, when one tries to generalize to further instances, whether or not one did so correctly — without being told just why one's attempt to generalize was correct or incorrect.

I understand if one is hesitant to grant that the student who is recognizing phenomenal qualities via explicit comparisons and deliberately applied mnemonic devices really possesses phenomenal knowledge of the quality in question. However, even if we decide to withhold an attribution of phenomenal knowledge in such cases, it becomes much harder to withhold after the student has stopped using the training device and is now able to recognize the property in question quickly and naturally, without any intervening artifice or technique.

Maybe this is a more reasonable view of how recognitional abilities are enhanced by instruction and practice. And maybe it's also a reasonable view of how we “naturally” — that is, without explicit instruction — come to recognize items in the world as belonging to one or another kind.

If so, however, then the smooth deployment of type-demonstrative concepts would still require a significant amount of knowledge on the part of the subjects who deploy them. This wouldn't be explicit knowledge of the salient features of the items recognized as “one of that kind,” or even of the similarities and differences that shaped the recognitional abilities in question.

Nonetheless, possession of these recognitional abilities — and thus the (type-demonstrative) phenomenal concepts they allow to be possessed — brings a lot of explanatory power to the account. It brings enough, for instance, to explain how Mary, in acquiring such dispositions after leaving her black-and-white room, manages to know so much.

So the physicalist may be able to have it both ways: phenomenal concepts can have the referential role of type-demonstratives, but their acquisition may still require a person to know a lot about the interrelations among the properties they denote. Awareness of an experience needn't provide immediate knowledge of what it's like, and knowledge of what it's like needn't involve knowledge of essence. The thesis of Revelation thus is unmotivated, which is a good thing, as it suggests a tighter than causal relationship between our experiences and our thoughts about them. As we have seen earlier, there is no such relationship for the having. Fortunately, none is needed.

3.5 Postscript: Imaginative Abilities and Phenomenal Representation

I remarked earlier that on one reading the quotational account of phenomenal concepts is functionally indistinguishable from an orthodox recognitional account that is supplemented by an additional requirement on phenomenal concept possession — namely, that a subject who possesses a phenomenal concept of Q must be able to *imagine* Q while tokening the phenomenal concept of Q. Presumably, such a requirement can be cashed out in terms of a subject's ability to token an *image* of Q, where this image is taken to be some phenomenal property representative of Q. The trick will be to come up with a plausible account of what it could mean for one phenomenal property (or property token) to be representative

of another. In this section, I explore what such an account might look like.

On the type-demonstrative view, possession of a phenomenal concept typically is linked to abilities to form images, which are used in the absence of currently perceived instances of a phenomenal property to demonstrate the property via deferred ostension. But such abilities are not, strictly speaking, required on type-demonstrative accounts such as Loar's. We might ask whether and how such abilities might be made a part of a recognitional account of phenomenal concepts.

I call this section a "postscript" because I am willing to accept the orthodox line of recognitional concepts, and hold, as discussed previously in this chapter, that the substantive nature of phenomenal knowledge can be explained by the fact that in at least some (and maybe most) cases, one has to know a lot about the properties in question, in particular their interrelations with others in the relevant quality space, to acquire the concepts.

Yet I sympathize with those who want more, and so this section explores a relatively innocent addition to the recognitional account that adds certain imaginative abilities as requirements on phenomenal concept possession. I will not pursue this line beyond this section.

I will consider a very general objection to purely recognitional accounts of phenomenal concepts — one which is milder than those previously considered and which may be met without occurring implausible conceptual or metaphysical baggage. The objection concerns the central role that imaginative abilities seem to play in phenomenal knowledge.

It may seem hard to conceive that one can know what it is like to have an experience of say, red, without being able to "imagine" red, in some sense when one's eyes are closed, or more generally when not actually having a visual experience of red. Quotational accounts have a ready answer: it is built into quotational accounts that one, in conceiving a phenomenal property, can token

or “echo” the property in question. But what about recognitional accounts? Can they account for the apparently obvious fact that in coming to know what it’s like to see red, I thereby come to be able to imagine red? Should they?

When we close our eyes and consider the difference between experiences of red and green, we are doing more than merely thinking of a distinction. Rather, it seems we are entertaining and contrasting two rich and distinct imaginings that differ in phenomenally substantive ways. This cannot be explained by reference to mere recognitional capacities. Dispositions to recognize contribute nothing to imaginative acts.

This worry, that demonstrative and recognitional accounts of phenomenal concepts do not capture the substantive nature of phenomenal thought, is a common one. Levine (2006) writes, “Our phenomenal concepts possess a substantive and determinate content; when entertaining phenomenal properties through them we have a conception of the specific phenomenal property, what, say, seeing red is like, as opposed to what seeing green is like. Demonstratives and indexicals are, with respect to substantive content, empty.”⁸

There are other reasons to think that recognition doesn’t tell the whole story. A concept may both be recognitional and refer to a phenomenal quality, and yet not be a phenomenal concept. We can imagine a neural hookup whereby subject A is able to reliably and directly recognize subject B’s tokening of phenomenal quality Q, without tokening it herself (nor any other phenomenal quality that co-varies with Q). In such a case, it seems that A has a recognitional concept of Q, but not a phenomenal concept of Q. We could also imagine various scenarios in which B alone is wired in such a way that when B tokens Q, B reliably forms the belief that Q is being tokened without consciously noticing Q in introspection.

Or we can imagine Jackson’s scientifically omniscient Mary, in her black and

⁸(Levine, 2006), p. 152

white cell, calculating exactly what physically introspectable effects seeing red would have on her, and thus developing a keen recognitional capacity. Still, no matter how evolved her skill, we would want to withhold from her the phenomenal concept, as she (most likely) lacks the ability to imaginatively represent red in her experience.

Earlier in this chapter, I insisted that when I think of, say, a shade of red, I do not thereby have an experience of red. Still, I seem to token *some* phenomenal quality that represents the shade in my experience. And this phenomenal representation is markedly different from the one I token when I think of green.

In light of these considerations, I propose the following additional requirement on phenomenal concept deployment:

Imaginative Constraint. *In order for a subject S to token a phenomenal concept of phenomenal quality Q , it is (conceptually or metaphysically) necessary that S be poised to token a phenomenal quality (or complex of qualities) Q^* , which phenomenally represents Q in experience.*

The claim that Q^* phenomenally represents Q in experience is actually two claims: 1) that one experiences Q^* *as representing*, and 2) that Q^* represents Q . I take the latter claim to be exactly where the dispositional-recognitional account comes in — i.e. the referential link between concept and quality is determined by the subject’s recognitional capacities. The qualitative nature of Q^* plays no criterial role in determining what property Q^* represents.

The notion of phenomenally representing is tied closely to what has been called ‘phenomenal intentionality’ (Loar, 2003). Phenomenal intentionality is, supposedly, an irreducibly phenomenal way of experiencing or imagining. When, for example, we see an object before us, we take the visual phenomenal qualities experienced as representative of something other than the qualities themselves. In fact, to imagine them *not* as representing requires effort. To sketch some seen

object, for example, requires taking the phenomenal qualities of the associated visual experience not as representative of an external object but as composing a visual field — a bit of phenomenal acrobatics that proves quite difficult for poor artists such as me.

It would be nice if there were a necessary relation between a phenomenal quality and the quality or qualities that represent it in experience. Unfortunately, there doesn't seem to be one. The relationship is complex, contingent, and highly variable. Consider the deployment of phenomenal concepts of determinables, such as the phenomenal concept of color. A

According to the imaginative constraint, I must be poised to imaginatively represent color in experience in order to wield the concept, and yet it seems I have an almost limitless supply of phenomenal qualities that will do the trick. This is unsurprising once we move away from the unclear idea that the “quasi-image” of Q tokened in phenomenal imagination or that the memory of Q is a “faint echo” of Q , as Tye (2003) puts it.

One can imaginatively represent Q via some Q^* without possessing the phenomenal concept of Q^* . For there is no guarantee that one can reliably recognize Q^* as Q^* in introspection, and thus it may not be the referent of any recognitional concept. This is good, for if possessing any phenomenal concept entailed possessing another, we would have circular dependencies or a vicious regress. But it does not, and we do not.

Chapter 4

Phenomenal Empiricism

In this chapter I consider a central point of agreement between the leading dualists' and the leading physicalists' accounts of phenomenal concepts. They all have as a consequence that possession of a phenomenal concept requires having had the phenomenal property that the concept denotes. I find reasons to reject this claim and describe a recognitional account of phenomenal concepts that does entail it.

4.1 The Thesis of Phenomenal Empiricism

The issue under consideration is whether possession of a phenomenal concept indeed requires having had the phenomenal property in question. An apparent difficulty in resolving this issue is that the term 'phenomenal concept' is a philosophical term of art.

We have no common conception of phenomenal concepts to rely on — no intuitive concept or term in natural language — in order to resolve the issue. This difficulty, however, is merely apparent, as the various accounts of phenomenal concepts share an aim.

The aim is to characterize a kind of concept that (a) denotes a phenomenal property and that (b) must be possessed by a subject who knows what it is like to have the phenomenal property in question. When the accounts are so viewed, we can ask whether the concepts that are characterized are phenomenal concepts,

that is, whether they in fact play this two-fold role.

In all the leading accounts under consideration, it is stipulated that the concepts denote a phenomenal property. And in all these accounts, the possession conditions for the concepts include having had the phenomenal property in question. Now it is indeed the case that knowing what it is like to have a phenomenal property requires having a concept that denotes the property. There is, however, an issue whether knowing what it is like to have a phenomenal property requires having had the phenomenal property.

Since on all the leading accounts having a phenomenal concept of a phenomenal property requires having had the phenomenal property, the issue will be whether knowing what it is like to have the phenomenal property requires having a phenomenal concept of it. If it doesn't, then the accounts fail to be correct accounts of phenomenal concepts, where such concepts are understood to be those concepts, whatever they are, that play the two-fold role mentioned above.

I will call the thesis that knowing what it is like to have a phenomenal property requires having had the property the thesis of *Phenomenal Empiricism*. More precisely, let Phenomenal Empiricism be the following:

Phenomenal Empiricism. *Necessarily, for any subject S and any phenomenal property Q, if S knows what it is like to have Q, it is at least partly in virtue of one or more experiences of Q.*

One of the goals of this chapter is to argue that Phenomenal Empiricism is false. We can know what it is like to have phenomenal properties we've never had. If that is indeed the case, then given that the accounts of phenomenal concepts in question aim to describe a kind of concept that denotes a phenomenal property and must be possessed by a subject who knows what it is like to have that property, they are incorrect accounts of phenomenal concepts.

4.2 The Leading Theories Entail Phenomenal Empiricism

As discussed in the previous two chapters, the leading theories of phenomenal concepts fall under one of two broad types: (i) type-demonstrative accounts, and, (ii) quotational accounts.

Not every theory can be neatly labeled as a type-demonstrative or quotational theory. For example: McLaughlin and Hill (1999) suggest that phenomenal concepts are used to recognize experiences sympathetically, rather than perceptually, Perry (2001) holds that the content of a phenomenal concept is reflexive, in that it depends always on external factors that locate the thinker relative to the subject matter, and Levine (2001) suggests that phenomenal concepts are non-ascriptive. These theories are not obviously committed to Phenomenal Empiricism, though they do not explicitly reject it.

According to type-demonstrative accounts, a phenomenal concept is a primitive mental representation that demonstrates a phenomenal property. We will focus here on the type-demonstrative account given by Loar (1990/1997), who holds that thoughts involving phenomenal concepts typically have the form ‘x is one of *that* kind’ or ‘x is another one of *those*’, where x is an occurrent experience and ‘that’ or ‘those’ expresses a phenomenal concept that picks out a phenomenal property of which x is an instance (Dualists also appeal to demonstrative phenomenal concepts. Chalmers (2003, p. 222), for instance, suggests that we employ demonstrative phenomenal concepts to pick out phenomenal properties in introspective contexts).

Type-demonstrative accounts entail Phenomenal Empiricism, typically by stipulating that phenomenal concepts always refer via events that tie them causally to instances of their referents. When a demonstration occurs in the context of an image, the image is held to be of a given kind at least partly in virtue of prior

experience of that kind. Tye (2003), for instance, suggests that we acquire introspective knowledge of what it is like to have an experience via a reliable process whereby a Q-experience triggers the application of a phenomenal concept of Q.

The second leading type of theory of phenomenal concepts is the quotational account (discussed in chapter three), which holds that a sample of a phenomenal property is actually a part of an instance of a phenomenal concept of that property.

Quotational accounts of phenomenal concepts also entail Phenomenal Empiricism. Although having a concept with a phenomenal property as a constituent does not entail having the phenomenal property in question, quotational accounts hold that a phenomenal concept is possessed partly in virtue of having had the phenomenal property that is its referent. The accounts also posit a central role for such phenomenal concepts in knowing what it is like.

Chalmers (2003, p. 223), for instance, holds that to know what it is like to have a phenomenal property is to know something along the lines of ‘C = R’, where R expresses a quotational phenomenal concept and C expresses some other sort of concept that picks out the phenomenal property rigidly, such as a demonstrative phenomenal concept.

4.3 Why Phenomenal Empiricism is False

4.3.1 Counterexamples

Phenomenal Empiricism is the thesis that knowing what it is like to have a phenomenal property requires that one have had it. As it stands, there seem to be obvious counterexamples. If I know what it is like to have an experience of a specific color shade, say, it seems I also know (or can come to know without further color experience) what it is like to have an experience of a slightly darker

shade, even if I have never had an experience of that shade. (Of course this sort of counterexample to empiricism was noticed by Hume, who dismissed it as “so particular and singular, that ’tis scarce worth our observing, and does not merit that for it alone we should alter our general maxim.”)¹

Similar examples exist in other sensory modalities: I can know what it is like to have an experience of a sweeter taste, a higher pitched tone, or a higher pressure touch than I’ve had previously. While I can know what it is like to have an experience of the slightly darker shade without having had such an experience, I cannot, according to the leading accounts of phenomenal concepts, possess a phenomenal concept of it. Thus, it seems the leading accounts have failed to characterize a kind of concept that is required for knowing what it is like to have a phenomenal property.

The proponents of the leading accounts might say that I don’t *really* know what it is like to have an experience of the darker shade or that the sense in which I know what it is like is not the sense at issue. It is true, of course, that I have some relatively uninteresting descriptive knowledge of the shade and that this kind of knowledge certainly isn’t what’s at issue. However, it seems that my knowledge of what it is like to have an experience of the darker shade goes beyond this.

Suppose, for instance, that I am looking at side-by-side color tiles of equally light shades, one of red and one of blue. Now suppose that a shadow is cast over the red one. It seems I can easily imagine how the blue tile would appear under the same shadow by attending to the hue of the blue tile and the lightness of the red tile. Once I am able to form an image of the darker shade (I leave it to section 3.2 to defend the claim that such an image would be *of* the darker shade), I no longer need to conceive of it in relation to the experienced shade.

¹(Hume, 2003), p.4

Rather, it seems that I am able to conceive of it in exactly the same way I can conceive of a previously had phenomenal property. I can, for instance, demonstrate it through deferred ostension, and come to be able to recognize it were I to have it.² (In deferred ostension, we point to one thing and in so doing refer to something else. As Quine (1977, p. 40) notes, deferred ostension comes naturally when we have a correspondence in mind. In the case at hand, the correspondence is between a type of image and the type of experience the image stands for.)

Two features of this sort of counterexample stand out: (1) it involves experiences that are phenomenally complex and (2) it involves phenomenally similar prior experience. It may be reasonable to suppose that our phenomenal knowledge of these unhad types of experience is due to prior experience of their various phenomenal aspects.

Given the simplicity of the diagnosis, one might wonder whether the leading accounts of phenomenal concepts need to characterize, for *any* phenomenal property Q, a type of concept of Q the possession of which is required for knowing what it is like to have Q. Perhaps they need only be committed to:

Weak Phenomenal Empiricism. *Necessarily, for any subject S and any phenomenal property Q, if S knows what it is like to have Q, it is at least partly in virtue of having one or more experiences that (taken together) include every phenomenal dimension of Q.*

Weak Phenomenal Empiricism has a good deal more psychological plausibility than Phenomenal Empiricism. For example, a subject who has had experiences of various shades of orange or purple plausibly could thereby know what it is like

²I take ‘able to recognize’ to mean something like ‘able to apprehend an instance of some antecedently conceptualized type as being of that type’. In this sense, one needn’t have prior experience of what is being recognized.

to have an experience of red, but a subject who has never had an experience of any red, orange, or purple very likely would not know what it is like to have an experience of red.

However, the leading accounts imply not just the *psychological plausibility* of Phenomenal Empiricism, but its *logical necessity*. Recall that type-demonstrative accounts have it that phenomenal concepts refer via events that tie them causally to instances of their referents, and quotational accounts have it that phenomenal concepts include the phenomenal properties that are their referents partly in virtue of those properties having been had. But consider that one could simply inventory the effects that having an experience of Q would have on the subject, and note that it is at least possible that those effects could occur in the absence of the cause.

This is well-illustrated by Dennett's (2006) "Swamp Mary":

Swamp Mary: Just as standard Mary is about to be released from prison [...] a bolt of lightning rearranges her brain, putting it by Cosmic Coincidence into exactly the brain state she was just about to go into after first seeing a red rose. (She is left otherwise unharmed of course; this is a thought experiment.) So when, a few seconds later, she is released, and sees for the first time, a colored thing (that red rose), she says just what she would say on seeing her second or nth red rose. "Oh yeah, right, a red rose. Been there, done that."³

Swamp Mary has never had an experience of red, but her brain is physically indiscernible from the brain Mary would have had post-release. Mary would have had an experience of red, and surely would have known what it is like to have such an experience. To deny that Swamp Mary knows what it is like to have an experience of red, one would have to reject the thesis that it is nomologically necessary that knowing what it is like supervenes on brain states or other states of the nervous system.

³(Dennett, 2006), p. 24

A representationalist, who holds that phenomenal properties supervene on intentional properties, might not agree that Mary and Swamp Mary have the same phenomenal properties. Of the theorists considered in this paper, only Tye is a representationalist (though he puts the issue of representationalism aside in his (2003) account of phenomenal concepts. Regardless of how the analysis goes for Swamp Mary, Tye's account still must deal with the previously given counterexamples, as he has it an a priori requirement that one must have had an experience of Q to possess a phenomenal concept of Q. It is consistent with representationalism to reject the claim that this is required for having a phenomenal concept.

Given the implausibility of rejecting the thesis that it is nomologically necessary that knowing what it is like supervenes on brain states or other states of the nervous system, it seems Swamp Mary can have the relevant thoughts about what it is like to have an experience of red and even weak Phenomenal Empiricism is refuted.

Granted, the content of some concepts is determined by causal-historical factors external to the subject, and thus Swamp Mary's conceptual stock may be impoverished or indeterminate in some ways. We can put the issue aside, however, as Swamp Mary's knowing what it is like does not require such concepts.

Furthermore, quantum theory holds that molecular re-arrangement cases like Swamp Mary are possible (though vastly improbable). It thus seems that counterexamples to weak Phenomenal Empiricism are not just logically, but nomologically possible.

I leave it an open empirical question whether there are counterexamples to weak Phenomenal Empiricism that apply to ordinary human beings in typical circumstances. Much work needs to be done on the phenomenology of various sensory modalities to adjudicate this. It is unclear, for example, what is going

on phenomenologically when we hear a tone and imagine it at a slightly higher pitch. Auditory illusions such as the Shepard scale and the tritone paradox might lead one to doubt that the tone one is imagining contains any simple phenomenal property that has not been previously experienced.

There is, perhaps, a way to preserve the leading accounts of phenomenal concepts in the face of the falsity of Phenomenal Empiricism. One could hold that knowing what it is like to have an experience of Q consists in having an *ability* or *disposition* to think certain kinds of thoughts about Q were one to have an experience of Q, where such thoughts involve phenomenal concepts (perhaps including thoughts like that expressed by Swamp Mary's "Been *there*, done *that*"). When Swamp Mary thinks such thoughts, she will be having an experience of Q and thus there is no difficulty explaining how she possesses the relevant phenomenal concept or concepts.

However, the leading accounts would still have failed in their original aim of specifying a kind of concept the possession of which is required for knowing what it is like, because before she has an experience of red, Swamp Mary knows what it is like to have it (in virtue of her abilities or dispositions) but fails to possess the relevant phenomenal concept.

Furthermore, there is the more general worry that one could **know that p** without being able to **think that p** (except in highly specific experiential circumstances). This seems highly implausible, and so I will not consider the option further.

If the quotational account were plausible, then Swamp Mary would not be a counter-example to Phenomenal Empiricism since she wouldn't be able to think about what it's like to see red prior to having an experience of red. But the quotational account is not plausible, or so I have argued (see Chapter Two), and Swamp Mary stands as a counter-example to Phenomenal Empiricism.

4.3.2 The Semantics of Phenomenal Concepts

In the face of the falsity of Phenomenal Empiricism, we considered and rejected abandoning the requirement that a theory of phenomenal concepts characterize a kind of concept of Q that must be possessed in order to know what it is like to have Q. The better course, I suggest, is to develop an account of phenomenal concepts in such a way that it doesn't require that a subject must have had Q to possess a phenomenal concept of Q, and therefore doesn't imply Phenomenal Empiricism. In this section, I will examine how such a modification might go.

The type-demonstrative account stipulates that a phenomenal concept of Q refers to Q partly in virtue of a subject's having had Q. I suggest that we simply do away with this stipulation, though some work needs to be done to explain what, on this modified account, makes a phenomenal concept of Q a concept of Q and not the phenomenal concept of some other phenomenal property or of no property at all. One way to go is with a naturalistic causal covariance account such as the following:

1. *A concept type C is the phenomenal concept of property Q if and only if C is a primitive concept and is the concept that the subject is disposed to token in various normal acts of introspection when Q is tokened, and because Q is tokened (i.e., no other concept screens C off from Q).*⁴

This naturalistic account of reference for phenomenal concepts does not require a subject have had Q to have a phenomenal concept of Q, as it takes a phenomenal concept of Q to refer to Q in virtue of a subject's disposition to token that concept when and because the subject has Q.

Alternatively, if one thinks that causal covariance accounts of representation

⁴*Roughly, A screens off B from C just in case all the correlation between B and C is in virtue of their correlation with A. Tye (2003) provides a good analysis of issues that arise for causal covariation accounts of representation.*

don't work (due to concerns about referential indeterminacy), one could be a deflationist about conceptual reference and stipulate that the phenomenal concept of Q refers to Q if it refers to anything. On this view, there are no empirical facts, causal or otherwise, that explain how C refers. Rather, the relation between C and Q is given by possession conditions for C that are analytic and a priori, and these possession conditions can require, for instance, that the subject be able to recognize Q when he or she has it.

One might fear referential indeterminism in cases where previously un-deployed phenomenal concepts are deployed for the first time through deferred ostension in imaginative thought. Such a worry would be based on the observation that an image used in deferred ostension could plausibly be an image of any of a variety of phenomenal properties.

The worry is unfounded, however, as the reference of a phenomenal concept deployed in imaginative demonstration is not fixed by any property of the image itself. Rather, its reference is fixed by the disposition of the subject to apply that same concept in introspective recognition (in the naturalistic case), or the reference is stipulated, and such dispositions are counted among the concept's possession conditions.

So Phenomenal Empiricism stands refuted, and an account of phenomenal concepts has been presented that does not imply it. This account preserves whatever virtues the type-demonstrative account has, while avoiding a mistaken commitment to Phenomenal Empiricism.

4.4 Purely Recognitional Phenomenal Concepts Vs. Phenomenal Demonstratives

First, a brief anecdote:

*The apocalypse is beginning, or so I was told recently by a man in Union Square. I should open my eyes and **recognize** it, he said. I replied that it was impossible to **recognize** this, and that I had an argument to prove it. Recognition, I said, was the action or fact of perceiving that some thing is the same as one previously known. Since it is an analytic truth that there can be only one beginning of the apocalypse, it cannot be recognized.*

Well, the street prophet paused, looked at me sadly, and handed me the heavy tome he was carrying. Lo and behold, it was not a bible but an abridged Oxford English Dictionary. The second definition of ‘recognition’ was highlighted in red. Recognition: the action or fact of apprehending a thing under a particular category, or as having a certain character. As my argument vanished in a puff of reasonable disambiguation, I slunk away, ears at the ready for the clop clop clop clop of a tetrad of horses...

Let us understand ‘recognition’ as the street prophet did: i.e., the action or fact of apprehending a thing under a particular category, or as having a certain character. Since just about all who take phenomenal concepts to be recognitional concepts endorse Phenomenal Empiricism, we should consider some of their reasons.

Loar doesn’t explicitly endorse Phenomenal Empiricism, but it follows readily from his claim that recognitional concepts are type-demonstratives. For it is a conceptual truth that a demonstrative cannot refer to a type that has not been tokened.

Phenomenal concepts are not demonstratives, or so I argue here and in chapter seven, but I admit there is a strong inclination to think they are. I think this feeling can be explained as a confusion between phenomenal concepts and other

demonstrative concepts that can be used to pick out phenomenal qualities, but which have little to do with “knowing what it’s like” (see chapter two) and thus do not fulfill one of the requirements on phenomenal concepts, i.e., that they be the sorts of concepts of phenomenal properties that one must possess in order to know what it’s like to have the properties in question.

If a phenomenal quality *Q* is present in experience, one may be able to wield a demonstrative concept of *Q*. For example, while staring at a color tile at time *t*, I am able to token a demonstrative concept *D* that has as its referent the specific phenomenal color quality *Q* that I am tokening in experience. Now, a central feature of demonstratives is that they change reference depending on context. So I cannot simply token *D* the next day and thereby refer to *Q*. I can, however, refer to *Q* descriptively as the referent of *D* at time *t*.

Now, suppose that it is the next day and I am tokening a phenomenal concept *C* of a particular shade of color. If I am deploying *C* with the hope that it refers to *Q*, as I am likely to do if my aim is to remember *Q*, I can consider whether (or hope that) the referent of *C* is identical to the referent *D* had at time *t*. Though I am tokening a demonstrative concept in my thought, I am not then engaged in an act of demonstration. I am mentioning the demonstrative, but not engaging in an act of demonstration.

Of course, a proponent of the demonstrative view of phenomenal concepts does not hold that these sorts of concepts regularly switch reference depending on context. This is because the emphasis of the account is on the idea of recognition, which, I suggest, we can more readily explicate without recourse to demonstration.

The main merit of the demonstrative account for a physicalist is that it honors a crucial requirement on any adequate theory of reference for phenomenal concepts: the concepts must refer “directly”, in the sense that the condition that determines reference needn’t be introspectively known, conceptualized, or even

conceptualizable. In other words, phenomenal concepts must be *representationally primitive*.

There are good reasons for the directness requirement. One is the regress threatened by the property dualism argument (see chapter seven). A theory which has me conceptualize qualities by way of some other conceptualized mode of presentation needs to explain how I come to conceptualize the mode of presentation. In the case of phenomenal qualities, it seems the only choices are a vicious regress of qualitative modes of presentation, or an embrace of analytic functionalism, neither of which appeals to an a posteriori physicalist.

Another reason to insist on directness or representational primitiveness is that it allows phenomenal concepts to be conceptually isolated from other sorts of concepts. One can wield all the non-phenomenal concepts in the world, but that is no guarantee that the conditions for the deployment of a phenomenal concept will be met. So, for example, there is no reason to believe that deployment of physical concepts will yield phenomenal concepts — or vice versa. This claim is vital to the a posteriori physicalist, who uses it to challenge the conclusions of knowledge and conceivability arguments.

Demonstrative accounts can be suitably direct, but they are ruled out by our denial of Phenomenal Empiricism. Still, the situation is not grim, as we can take a causal recognitional account and recast it in dispositional terms.

According to Tye, “Phenomenal concepts refer via the causal connection they have with their referents. In first approximation, a phenomenal 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.”⁵

⁵(Tye, 2003), p. 7

We can avoid committing to Phenomenal Empiricism by casting Tye's requirement in dispositional terms: In first approximation,⁶ a phenomenal concept C refers to a phenomenal quality Q via C's being the concept that *would be* exercised in an introspective act of awareness by person P if, and only if, under normal conditions of introspection, Q were tokened in P's current experience and because Q were tokened.

To determine the referent of a phenomenal concept, we ask which phenomenal quality if any, would, under normal conditions of introspection, cause that concept to be tokened. As the relevant dispositional connection is merely causal, we get the directness we need, and as it is merely dispositional, we can deny Phenomenal Empiricism.

⁶There are general issues surrounding causal covariation accounts of representation that would need to be addressed for a full account. Tye (p. 7), provides a good analysis of the issues.

Chapter 5

Recognitional Concepts: Reply to Fodor

It is my intention in this work to keep the focus on phenomenal concepts and not to get mired in foundational debates over the nature of concepts generally and their role in a general theory of cognition. For the most part, how such foundational debates shake out does not affect the present project (barring the most skeptical of conclusions), except perhaps to provide an interpretation for it. However, it seems appropriate to move a branch lower in the dialectical tree to take on a recent argument by Jerry Fodor, which purports to show that there are no recognitional concepts.

I take this on here for three reasons. First, because the theory of phenomenal concepts I defend requires that there be recognitional concepts; second, because the argument I present against Fodor may not generalize to a defense of all recognitional concepts, and thus may be of interest specifically to those interested in phenomenal concepts; third, because I think Fodor's attack brings out an important distinction that needs to be made between a nominalist account of recognitional concepts and their referents and a realist account. I will argue that Fodor's argument addresses the former sort, but leaves the latter unscathed, and that phenomenal concepts are best understood as recognitional concepts in the realist vein.

5.1 Fodor's Argument

Following Fodor, let us start with the following definition:

Recognitional. *A concept is recognitional if and only if among its possession conditions is the ability to recognize at least some things that fall under C as falling under C.*

Let us agree that concepts are partially constituted by their possession conditions; that is, that they have their possession conditions essentially. It follows then that C is a recognitional concept if and only if it numbers, among its possession conditions, the ability to recognize at least some of the extension of C as falling under C.¹

Key to Fodor's argument is a familiar compositionality constraint:

Compositionality. *S is a satisfier for concept C if and only if C inherits S from the satisfiers for its constituents concepts.*

What makes Compositionality plausible is a well-known, albeit controversial, story about the productivity of natural language and constraints on how natural languages are learned. Compositionality forms the basis of an explanation of why concepts are systematic and productive; and, *mutatis mutandis*, of how it is possible to learn a language by learning its finite basis. Let us grant Fodor the principle for the present, though we will revisit it a bit later.

Fodor claims that Compositionality is violated if there are complex recognitional concepts with recognitional concepts as constituents. To see this, he asks us to note that that one could be in position to recognize good instances of a type

¹Fodor appears to use the terms 'ability' and 'capacity' interchangeably - both to indicate an actualized recognitional talent. I will use only 'ability' to mean this, reserving 'capacity' to indicate something like an unpossessed ability that it is possible (in some nomological sense of possibility) for a subject to acquire. For example, most people have the capacity, but not the ability, to play the piano passably.

designated by a complex concept without being in position to recognize good instances of types designated by its constituent concepts. If, for example, PET, FISH, and PET-FISH² are recognitional, one plausibly could possess PET-FISH partly in virtue of the ability to recognize good examples of pet fish (goldfish, guppies, etc.), and yet not be able in general to recognize good examples of pets (dogs, cats, etc.) or fish (trout, sharks, etc.).

In this case, one would possess PET-FISH, but neither PET nor FISH, violating Compositionality. This, in a nutshell, is Fodor's case against recognitional concepts.

We might reconstruct Fodor's argument as follows:

- (1) If PET and FISH are recognitional concepts, then PET-FISH is a recognitional concept (no reason to suppose otherwise).
- (2) A subject might be able to recognize pets and fish without being able to recognize pet fish.
- (3) If PET and FISH are recognitional concepts, then a subject might possess PET and FISH without possessing PET-FISH.
- (4) A subject cannot possess PET and FISH without possessing PET-FISH (Compositionality).
- (5) Therefore, PET and FISH are not recognitional concepts.

The most common reply to Fodor is to deny (1). One might hold that recognitional abilities are possession conditions for some primitive concepts, but that the complex concepts formed from them have exactly those recognitional possession

²In this chapter, I will use small caps to represent concepts. Also, for simplicity, I will use the notation CONCEPT₁-CONCEPT₂ to generically indicate the joining of two concepts into a complex concept. Thus, the notation will be used both in cases where the joining is intersective (e.g. RED-APPLE), and those where it is not (e.g. IMITATION-GOLD).

conditions of their constituent concepts. In the example above, then, one needn't be able to recognize good examples of pet fish to possess PET-FISH, but one would need to be able to recognize good examples of pets and good examples of fish.

This is the solution given by Schiffer (1999).

In rejecting Schiffer's view, Fodor writes,

If one's grasp of PET-FISH isn't constituted by one's ability to recognize its instances, then, surely, one's grasp of FISH isn't constituted by one's ability to recognize its instances. And vice versa: if recognizing good instances like goldfish isn't a satisfier for the concept PET-FISH, then recognizing good instances like (as it might be) trout, isn't a satisfier for the concept FISH. Surely, being able to recognize goldfish as pet fish must stand in precisely the same relation to having the concept PET-FISH that being able to recognize trout as fish stands in to having the concept FISH. So, how could that relation be constitutive of concept possession in the one case but not in the other?³

I think, in general, Fodor's point is a good one, as far as pets and fish go (though he does rely quite heavily on the proof surrogate 'surely'). To have the concept PET-FISH you have to have the concept FISH. But why suppose that to have the concept PET-FISH you have to have a recognitional capacity for good instances of fish? The proposal seems unmotivated. Why should the ability to recognize, say, trout or kittens, play into the concept possession conditions for PET-FISH?

The answer, I think, is that recognitional abilities *do* compose, but not in the way Fodor thinks they should. Specifically, the ability to recognize *good* or *typical* cases does not compose, but, I shall argue, there is no reason to think that recognitional abilities need be tied to good or typical cases. Once this is cleared up, we can see that recognitional abilities do, in fact, compose in a way which is sufficient for the recognitional account of phenomenal concepts. The premise that recognitional abilities must always apply to good instances is plausible only on a nominalist metaphysical understanding of the nature of what it is that is

³(Fodor, 1998), p. 122

being recognized. Fodor's criticism, I suggest, is a criticism of a certain brand of nominalism a brand which a physicalist who holds a realist view of phenomenal properties should reject.

5.2 Nominalism, Realism, and Fodor's Argument

In this section, I wish to a) draw a distinction between nominalist and realist conceptions of recognitional concepts, b) show that Fodor assumes a nominalist conception of recognitional concepts, c) show that Fodor's argument has no force against a realist conception, and d) suggest that the realist conception is better for the physicalist to hold, at least as far as phenomenal concepts go.

5.2.1 Nominalist and Realist Conceptions of Recognitional Concepts

There are a couple of ways to understand what's going on referentially when one deploys a recognitional concept. The first way, staunchly in the nominalist tradition, holds that the denotation of a recognitional concept is *defined as* whatever it is that regularly triggers the tokening of the concept. On this view, the world is carved by our concepts, and our concepts successfully refer with a minimum of fuss and without any skeptical worries merely in virtue of their being applied in the way we apply them.

On a strictly nominalist view, the world is divided into elms and oaks, cats and dogs, odds and evens, the beautiful and the ugly ultimately in virtue of how we apply the concepts ELM and OAK, CAT and DOG, ODD and EVEN, BEAUTIFUL and UGLY. This is not to say that our concepts are brute or static. On the contrary, the nominalist is as capable as anyone of fine-tuning ontology as science progresses. The difference is that the nominalist ultimately will explain this

in terms of conceptual change and theory development rather than ontological discovery.

An alternative to the nominalist conception of recognitional concepts is the realist conception. On this view, the denotations of recognitional concepts are real properties (or substances) that are individuated independently of the concepts that refer to them. On the realist account, recognitional concepts refer successfully in virtue of picking out real properties — the concepts themselves do nothing to individuate the properties. On a strictly realist view, there is a distinction *in the world* between elms and oaks, cats and dogs, odd and even, the beautiful and the ugly — a distinction that has nothing to do with the way we conceive elms, oaks, cats, dogs, odd, even, the beautiful, the ugly, etc.

Referential success in concept deployment depends at least in part on the ability of a subject to “latch on” to a real property, and, in the case of recognitional concepts, to trigger a token of a concept in response to its referent. Of course, one needn’t choose sides here. One might, for example, choose to be a realist about basic physical properties and a nominalist about beauty. There are indefinitely many permutations and it is not my intention to stake out a position on this (perhaps the most foundational of?) foundational issues.

5.2.2 Fodor’s Assumption of a Nominalist Conception of Recognitional Concepts

Fodor writes:

For one thing, insofar as recognitional capacities are construed as perceptual capacities, the claim that there are recognitional concepts preserves the basic idea of Empiricism: that the content of at least some concepts is constituted, at least in part, by their connections to percepts [...] More generally, the claim that there are recognitional concepts is a bastion of last resort for philosophers who think that semantic facts are constituted by epistemological facts[...] for example, a familiar line of anti-skeptical

argument turns on there being some. The idea is that, if a concept is recognitional, then having certain kinds of experience would, in principle, show with the force of conceptual necessity that that the concept applies. If, for example, RED is a recognitional concept, then having certain kinds of experience would, in principle, show with the force of conceptual necessity that there are red things.⁴

It is clear from this passage that Fodor understands recognitional concepts as belonging to the nominalist's toolbox. After all, it is only within a nominalist perspective that a recognitional account of concepts could, by itself, form any kind of defense against skepticism. On a realist perspective, it is possible to *misapply* recognitional concepts. It would be possible, for instance, to mistakenly judge objects as red that aren't, or even experiences of various sorts as experiences of red, that aren't. On the realist view of phenomenal properties, one cannot "show with the force of conceptual necessity" that experiences of various types exist *merely* by consulting one's recognitional concepts.

Further evidence that Fodor is assuming a nominalist perspective comes from his frequent claim that recognitional concepts *must* apply, "with the force of conceptual necessity", to "good cases" or "good examples". On a nominalist conception of recognitional concepts, this is plausible. Vagueness is to be found at the edges of our concepts, not in the world, and so, by necessity, cases where recognitional concepts clearly apply necessarily coincide with "good" examples of what such concepts denote.

5.2.3 Compositionality on a Realist Conception of Recognitional Concepts

Suppose with the realist that recognitional concepts are concepts of real properties (or substances). They involve fallible abilities (not mere dispositions) to

⁴(Fodor, 1998), p. 128

successfully recognize these properties under a variety of conditions. In the ordinary case, some property P may be picked out by a multitude of concepts, some of which may be recognitional, some of which may be discursive. Even in the case of phenomenal properties, there may be multiple kinds of recognitional concepts which pick them out. Some of these are applied introspectively and directly — these we are calling ‘phenomenal concepts’.

But there may be other recognitional concepts of phenomenal properties. I may recognize a phenomenal property via behavioral observation or by attending to a correlated phenomenal or non-phenomenal property. Perhaps I could reliably point to a blip on a brain scan. The point is that once we grant that a phenomenal property is not *defined* by some concept that refers to it, we can see that the role that recognition plays is a referential role, pure and simple. It plays no part in defining its referent.

The upshot is that whether I can think about P has nothing (conceptually) to do with *how* I think about P. Once reference is achieved, however it is achieved, I can kick away the ladder, so to speak, and think about the referent. This is in stark contrast to the nominalist take, in which *how* I think about P is inextricable from *what P is*. This is manifested, in the case of recognitional concepts, by attending to good cases, or stereotypes. Fodor is right that stereotypes do not compose, and, if recognitional concepts were individuated by the stereotypes they pick out, neither would they. But on the realist conception, this is not how concepts are individuated. They do not individuate their referents. Rather, they are, we may suppose, individuated essentially *by* their referents.

It may have been better if the term ‘recognitional concept’ were replaced by ‘recognitional conception’, so as to separate the manner in which a concept is applied, the conception, from the essential possession conditions of the concept (e.g., that it have the content it has). A ‘theory of phenomenal concepts’, then, would

become a ‘theory of phenomenological conception’ — a study of the special way we acquire concepts of phenomenal properties from a direct, introspective perspective. On this clearer terminology, the term ‘phenomenal concept’ would apply to all concepts of phenomenal properties, not just those from a direct, introspective perspective. I will, of course, stick to the common usage of ‘phenomenal concept’, though it be (possibly) the source of some confusion.

Let us bring the discussion back now, first to Fodor’s case of pet fish and then to phenomenal concepts. To have the concept FISH, one must be able to refer to the property *fish*. There are several routes to referring to fish, some discursive and some recognitional. Focusing solely on the latter, we may say that one has a recognitional concept FISH if and only if one is able to pick out fish. Now, a snorkeler and a gourmand, for instance, may both be able to recognize fish, one by sight and one by taste, and yet there is a clear sense in which they are able to converse on the same topic: i.e., fish. This is true even though their stereotypes (or ‘good examples’) are quite different. The snorkeler might look for a characteristic shape or movement, and the gourmand may look for a characteristic taste or texture. An eel may throw off the snorkeler but not the gourmand, and halibut vice versa.

Let us also assume, for the sake of exposition, a realist conception of PET. That is, let us assume there really are such things as pets in the world — and that this fact holds independently of peoples’ dispositions to classify various items as pets. Also, let us assume for the sake of simplicity that the only pet fish in existence are goldfish, and that the gourmand, used to large fillets and steaks, would not recognize a goldfish as a fish. Finally, let us assume that the gourmand has no other concept (or conception) of fish.

Does it follow that the gourmand does not have a concept PET-FISH? Of course not. The gourmand can think about pet fish even though he is unable to

recognize any (actual) pet fish. He may think to himself “I hear that some fish are also pets. I should like to taste one of these pet fish some day!” Compositionality for recognitional concepts holds even as certain recognitional abilities, namely abilities to recognize typical instances, do not. The gourmand would possess a concept PET-FISH, even though he couldn’t pick out goldfish.

It is worth noting that the recognitional abilities of the gourmand *do* compose — just not in a way that is of use in the actual world. For example, the gourmand would be very good at spotting pet fish in a world in which swordfish and salmon were kept as pets. Perhaps there is a pet swordfish in the actual world. It would not count as a good example, solely because of its atypical nature, but it would be recognizable to the gourmand.

We now arrive at the heart of the problem with Fodor’s argument. His core claims can be reconstructed as follows:

- (1) To possess a recognitional concept of X is to be able to recognize good instances of X.
- (2) A subject may be able to recognize good instances of P, good instances of Q, but not good instances of (P and Q).
- (3) A subject S may possess a recognitional concept of P, a recognitional concept of Q, but not a recognitional concept of (P and Q). (1, 2)
- (4) (Compositionality)
- (5) Therefore, there are no recognitional concepts. (3,4)

Schiffer wants to modify (4) by claiming that not all the possession conditions need compose. I suggest that the problem is with (1). The notion of a good instance of a property is necessarily related to our ability to identify that property only on a nominalist conception of that property. Once we become realists about

the property, we either reject the notion of a good instance altogether, or explicate it independently of our conceptual apparatus. Recognitional abilities are diverse. Each is better suited to some contexts than to others.

When our recognitional concepts compose, so do our recognitional abilities, though the contexts in which they apply may change. In some cases, this will result in an equal (or better!) ability to recognize the property conjunction, but in some contexts this ability will diminish or apparently vanish. This is not surprising in the least, however, as abilities can only be exercised (or acquired) in certain contexts. But one can still have the ability when contexts are changed.

This seems the natural way to look at things. Suppose that our snorkeler, who possessed a concept of fish in virtue of her ability to visually identify fish, suddenly goes blind. Such a blind snorkeler would still be able to think thoughts of fish, even as she has lost the ability to reidentify fish that allowed her to think about fish in the first place. What this shows is that there is an ambiguity in the notion of a ‘possession condition’ for a concept. There is a sense in which concepts are partially constituted by their possession conditions, and another sense in which they are not. They are not if ‘possession conditions’ is taken to include facts about how the subject picks out the referent.

Recognitional abilities, modes of presentation, sub-personal causal dispositions, and the like, can all be part of this story, but none of them are essential to being able to think about some real substance or property. But there is another, more basic sense, in which concepts *do* have their possession conditions essentially, and in which these possession conditions compose (and decompose): one possesses a concept of P if and only if one is able to think about P. Of course, this is exactly the position that Fodor holds and which he thinks is incompatible with the existence of recognitional concepts. But it is not. It is incompatible only with a nominalist account of recognitional concepts.

I opened this discussion by claiming to provide a limited defense of recognitional concepts, a defense which covers phenomenal concepts but which may leave others behind. There are myriad purportedly recognitional concepts which do not fit easily into a realist ontology. PET may be such an example. For those who wish to posit a range of recognitional concepts that are best understood nominalistically, Fodor's challenge stands. It is sufficient for this work, however, to defend recognitional *phenomenal* concepts, and these concepts, I maintain, are best understood on a realistic conception of phenomenal properties.

It should be clear from the bulk of this thesis that I am assuming a realist ontology of phenomenal properties. For example, I take it that one may have distinct experiences of unique red or red-31, regardless of whether one has the ability or disposition to distinguish among such experiences in thought. (This is not to say that, in fact, the range of experiences we can have is independent of our conceptual abilities. There may be complex causal interdependence in some cases).⁵ What makes a phenomenal concept of red a phenomenal concept of red is that it *in fact* picks out experiences of red. In no way does it *define* what an experience of red is.

It is beyond the scope of this thesis to fully defend this point, as it takes us away from the nature of phenomenal concepts to the nature of phenomenal properties themselves. Suffice it to say that the realist approach seems most congenial to the scientific tradition of bold, risk-laden, multifaceted analyses of natural phenomena. Close analysis, be it phenomenological, hetero-phenomenological, or neurological analysis, for example, could conceivably lead one to doubt that one is sorting one's experiences correctly according to phenomenological type. At the

⁵Such interdependencies may be more evident in emotional experiences than in perception. There may be, for instance, phenomenal properties associated with *dread*, that can only be had by creatures that are able to deploy concepts of the self, the future, and of danger. I will remain agnostic on this point.

least, a physicalist should hold open this possibility.

While phenomenological analyses may well prove to be important sources of evidence, they are only part of the puzzle. We may, in extreme cases, override a person's sincere avowal that she is having a particular kind of experience, say of pain, in the face of recalcitrant evidence of another sort. For the realist about phenomenal properties, such skepticism cannot be ruled out a priori.

5.3 Are There Complex Phenomenal Concepts?

The previous section sought to vindicate all those recognitional concepts that have, as their denotations, independently real properties or substances. In this section, I explore another way to confront Fodor's argument, but this line of reasoning does not generalize quite so well. I propose to defend the claim that there are no complex phenomenal concepts.

Let me start by making clear what I *don't* mean when I say there are no complex phenomenal concepts. First, I don't mean that there are no complex concepts that have phenomenal concepts as constituents. There are myriad ways in which phenomenal concepts can compose. Some complex concepts are composed of both phenomenal and non-phenomenal concepts, (e.g. YESTERDAY'S-PAIN, where 'pain' is understood phenomenologically). Such complex concepts clearly are not phenomenal concepts (nor recognitional concepts of any sort).

We may also form complex concepts solely from phenomenal concepts. For example, if I possess phenomenal concepts of two types of taste experiences, SWEET and SALTY, I can without difficulty form a complex concept SWEET-SALTY. What I intend to argue in this section is that insofar as concepts like SWEET-SALTY are complex, they are not phenomenal, and insofar as we can possess a phenomenal concept SWEETSALTY, it is not complex (though the phenomenal property it

denotes *is* phenomenally complex).

It is crucial here to note the distinction between a complex concept and a simple concept of a complex of properties. Only the former, I suggest, is never a phenomenal concept. To illustrate, consider three phenomenal properties: the particular experienced sounds of an occurrence of a piano playing the notes middle-C, E, and G. Call the three specific felt qualities of the respective notes Q^C , Q^E , and Q^G , and the respective phenomenal concepts of the qualities C, E, and G. Together, Q^C , Q^E , and Q^G form a specific experience of the chord C-major — call that experience Q^{CEG} .

Now, it is plausible that although one cannot experience Q^{CEG} without experiencing Q^C , Q^E , and Q^G , one nevertheless can recognize or think about Q^{CEG} without recognizing or thinking about Q^C , Q^E , or Q^G . Thus, one can possess the phenomenal concept CEG without possessing C, E, or G.⁶ Thus, it is reasonable to suppose that CEG is not a complex concept, but a simple phenomenal concept.

I suggest that all phenomenal concepts are simple, and that there is no sense to be made of a complex phenomenal concept. Note that this is not to deny the obvious truth that phenomenal concepts can be components of (non-phenomenal, non-recognitional) complex concepts. Rather, the point is just that the PET-FISH problem never arises for phenomenal concepts because there are no complex phenomenal concepts, and thus one can deny the existence of complex recognitional concepts without being ad hoc about it.

But what of truly complex concepts with phenomenal concepts as constituents? What are the possession conditions for the complex concept C-E-G, for example? The answer is that they are exactly the possession conditions for C, E, and G.

⁶Even more to the point, each experienced note comprises various experiences of overtones and undertones which may not be recognized or conceptualized.

This is plausible because we have an entirely different simple phenomenal concept, namely CEG, if we are interested in thinking about the experience of the conjunction of the notes, rather than the conjunction of the experiences of the notes.

To take a more familiar example from the literature, consider a specific type of color experience. It is widely accepted that one may possess a phenomenal concept of the experience type, even as one is unaware that the color experience is composed of phenomenally distinguishable attributes: hue, lightness, and saturation being the most obvious examples. A trained subject can distinguish these attributes in experience, but doing so is certainly not necessary for possessing a simple phenomenal concept of a type of color experience.

Let us now return to Fodor's argument. In many cases, one can detect good instances of Ps and good instances of Qs, but not good instances of (P and Q)s. Fodor took this to show that recognitional concepts violate compositionality. Schiffer suggested that possessing the complex concept P-Q requires only being able to recognize Ps and Qs, and not (P and Q)s. Fodor found this suggestion unmotivated. He writes:

Couldn't we just split the difference? Couldn't we just say that the satisfiers for the primitive concepts include recognitional capacities, even though the satisfiers for complex concepts don't? [...] Simply not credible. If one's grasp of PET-FISH isn't constituted by one's ability to recognize its instances, then, surely, one's grasp of FISH isn't constituted by one's ability to recognize its instances. And vice versa: if recognizing good instances like goldfish isn't a satisfier for the concept PET-FISH, then recognizing good instances like (as it might be) trout, isn't a satisfier for the concept FISH. Surely, being able to recognize goldfish as pet fish must stand in precisely the same relation to having the concept PET-FISH that being able to recognize trout as fish stands in to having the concept FISH. So, how could that relation be constitutive of concept possession in the one case but not in the other?⁷

⁷(Fodor, 1998), p. 131

Here, now, we have an answer, at least in the case of phenomenal concepts. If two phenomenal concepts are constituents of a complex concept, there is no reason to suppose that there is a conjunction of their referents which should be recognizable independently of the referents considered singly. In the PET-FISH case, the good examples of pet fish diverge greatly from the good examples of pets and the good examples of fish. But this is not the case for SWEET-SALTY. Insofar as there is an experience type (sweet and salty), we could think about it using a simple phenomenal concept SWEETSALTY. Unlike SWEETSALTY, there is no simple concept PETFISH, and therein lies the difference.

There are interesting phenomenological questions which will not be pursued here. For example, when I think of an experience of seeing a red triangle, do I token a complex non-phenomenal concept with phenomenal concepts as constituents, i.e. RED-TRIANGLE, or a simple phenomenal concept, i.e. REDTRIANGLE? It seems to me that the former is more plausible, as there seems an analytic entailment (in the classical sense) from the thought of a sensation of seeing a red triangle to the thought of a sensation of seeing a triangle. Contrast this with the case of the chord discussed above, where there is no such entailment from a thought about the sensation of hearing a chord to a thought about the sensation of hearing a constituent note. In general, I would suppose that concepts of sensations involving multiple modalities are complex and non-phenomenal, though cases of synesthesia may point to this being merely a contingent truth or general rule.

Chapter 6

The Phenomenal Concept Strategy: Reply to Stoljar

6.1 Chapter Introduction

The goals of this chapter are twofold. The first goal is to defend the claim that the physicalist's phenomenal concept strategy does not require the assumption of Phenomenal Empiricism (see chapter three). The second goal is related to the first. It is to defend the phenomenal concept strategy against criticisms recently put forward by Daniel Stoljar (2005), who claims 1) that the phenomenal concept strategy relies on the truth of something like Phenomenal Empiricism, and 2) that the strategy fails even if the thesis holds.

While I agree with Stoljar that the physicalist should not rely on theses such as Phenomenal Empiricism (in part because they are false, or so I would argue), I think Stoljar misunderstands the way the strategy has been (or at least should be) applied. I think Stoljar's arguments, while ultimately unsuccessful, are useful in that they underscore the irrelevance of theses like Phenomenal Empiricism to the phenomenal concept strategy. If the considerations in this chapter are correct, Phenomenal Empiricism is not only false, but even if it were true, it would be of no help to the physicalist.

It might seem that in rejecting Phenomenal Empiricism the physicalist concedes too much — that without that thesis (or something close it), the physicalist cannot hope to explain how purportedly necessary psychophysical identities could

possibly be a posteriori. Such a view is understandable, considering that much attention has been given to issues that may turn on the truth of Phenomenal Empiricism, such as whether scientifically omniscient Mary can know what it is like to have an experience of red prior to having such an experience. It is not surprising that some, such as Stoljar, have gone so far as to characterize the strategy of appealing to phenomenal concepts to answer antiphysicalist arguments as the strategy of appealing to theses like Phenomenal Empiricism.

6.2 Varieties of the Phenomenal Concept Strategy

Before addressing Stoljar's attacks on the phenomenal concept strategy, it is important to get clear on what he takes the strategy to be, and how that relates to the strategy defended in this thesis.

Stoljar provides the following definition:

Phenomenal Concept Strategy. Phenomenal concepts are importantly different from other concepts, and arguments against physicalism, in particular the knowledge and conceivability arguments, fatally neglect to take this difference into account.

There is a great deal at stake here. If the phenomenal concept strategy succeeds, it would show that many of the most prominent puzzles in the philosophy of mind are due to a failure to note fairly obvious facts about our conceptual framework. Success of the strategy would provide an attractive deflation of these puzzles, and would deflate them without insisting on any actual or potential progress in the sciences, nor on any counter-intuitive claims of analytic entailments from biological, psychofunctional, or behavioral facts to facts about felt experience.

A variety of authors have espoused some version of the phenomenal concept

strategy. Loar, in perhaps the most famous articulation of it, suggests that understanding the unique nature of phenomenal concepts will allow us to “take the phenomenological intuition at face value, accepting introspective [i.e. phenomenal] concepts and their conceptual irreducibility, and at the same time take phenomenal qualities to be identical with physical-functional properties of the sort envisaged by contemporary brain science”.¹ Tye writes that the apparent force of the knowledge and conceivability arguments stems “largely from a failure to recognize the special features of phenomenal concepts”.²³ Hill and McLaughlin write “because of certain facts about sensory [i.e. phenomenal] and physical concepts” the arguments can be answered.”⁴ Other notable philosophers employing some version of the phenomenal concept strategy include Block and Stalnaker (1999), Perry (2001), Sturgeon (2001), Papineau (2001), and Block (2006).

There are a number of ways the phenomenal concept strategy can proceed, but let us follow Stoljar in assuming that the central move in the phenomenal concept strategy is to focus on the *possession conditions* for phenomenal concepts.

Not all versions of the phenomenal concept strategy focus on concept possession conditions. There are those which focus rather on concept *acquisition* conditions, such as Papineau (1999). For Sturgeon (2001), the emphasis is on concept *application*, where a thesis about concept application is an epistemological thesis about what conditions must be met in order that one is justified in applying a phenomenal concept in the first place. For Hill and McLaughlin (1999), the basic idea is developed via a discussion of the epistemological norms

¹(Loar, 1990/97), p. 598

²(Tye, 2003), p. 707

³Tye no longer thinks the phenomenal concept strategy works. I will not examine his new position in this thesis, as the publication in which he espouses his new view is forthcoming as of this writing.

⁴(Hill and McLaughlin, 1999), p. 447

that govern the *faculties* that are involved in the deployment of the relevant concepts. For them, what is important is the fact that a phenomenal concept bear a constitutive relation to introspection and judgment.

Stoljar focuses on the following thesis, the truth of which he believes the phenomenal concept strategy is designed to establish:

Experience Thesis. S possesses the (phenomenal) concept C of experience E only if S has actually had experience E.

Stoljar writes, “the suggestion of the phenomenal concept strategy is that the experience thesis entails that phenomenal concepts are very different from other concepts, and this difference can be exploited to disarm the arguments against physicalism.”⁵

It is easy to see why the Experience Thesis would appeal to a physicalist. If it were true, it would clearly demarcate an epistemically striking and seemingly important feature of phenomenal concepts. For non-phenomenal concepts, there is no sort of experience that would strike one as necessary for possessing that concept. But in the case of phenomenal concepts, one might seek to identify the very referent of the concept as filling the role of a necessary possession condition of the concept itself.

Stoljar is willing to concede that not all versions of the phenomenal strategy require the Experience Thesis to be true. He thinks Loar might not endorse it, for instance. Based on conversation with Loar, I think, to the contrary, that he does endorse the Experience Thesis. But that is unimportant for present purposes. What matters is that Stoljar claims that his case against the phenomenal concept strategy proceeds just as well against a version of the strategy which denies the Experience Thesis, but relies on the following alternative claim:

⁵(Stoljar, 2005), p. 471

Recognition Thesis. S possesses the (phenomenal) concept C of experience E only if S has certain dispositions to recognize, discriminate and identify E if S has or undergoes E.

The difference between the Recognition Thesis and the Experience Thesis is that only the latter implies that having the experience is necessary for having the concept. Given the considerations against Phenomenal Empiricism advanced in chapter three, I would, of course, endorse the Recognition Thesis while denying the Experience Thesis.

Stoljar does not provide a separate argument against versions of the phenomenal concept strategy that rely on the truth of the Recognition Thesis, but it would seem that one is required, insofar as I am correct that his case against the phenomenal concept strategy fails.

Unlike me, Stoljar does not reject the Experience Thesis. Rather, he argues that the thesis cannot be used successfully to reply to the antiphysicalist arguments. I do not dispute this; however, I do dispute his further claim that this shows that phenomenal concept strategy has no chance of success.

Stoljar offers two arguments to support this latter claim, but these arguments fail, or so I shall argue. His first argument moves from the observation that there are certain kinds of a priori connections between phenomenal concepts and other concepts to the conclusion that the phenomenal concept strategy cannot be used to explain how a statement that expresses the physicalist claim that the physical facts entail the phenomenal facts could be both necessary and a posteriori. His second argument turns on the claim that if the phenomenal concept strategy were successful against the conceivability argument against physicalism, then it would be successful against a similar, but obviously sound, conceivability argument against behaviorism. I will examine these arguments in turn.

6.3 Stoljar's First Argument

Let P be a conjunction of the physical truths of our world and P^* be a conjunction of the phenomenal truths. We may say, following Jackson (1998), that physicalism entails the necessary truth of the psychophysical conditional: 'If P then P^* '. (I am unsure that the physical truths are few enough to be put in a statement, even in one of infinite length. But I will follow Stoljar and Jackson's characterization here, as nothing turns on it for present purposes.)

Stoljar distinguishes an a priori conditional from what he calls an "a priori synthesizable" conditional (p. 478):

A Priori. 'If A then B ' is a priori if and only if a sufficiently logically acute person who possessed only the concepts required to understand it, is in position to know that it is true.

A Priori Synthesizable. 'If A then B ' is a priori synthesizable if and only if a sufficiently logically acute person who possessed only the concepts required to understand its *antecedent*, is in a position to know that it is true.

On Stoljar's view, if Phenomenal Empiricism were true, we would have an explanation of why the psychophysical conditional isn't a priori synthesizable, but not of why it isn't a priori. After all, there are many truths that are a priori but not a priori synthesizable. Stoljar gives as an example:

6.3.1. *If x is rectangular, then x has some property or other.*

Presumably, Stoljar claims, one might possess the concept RECTANGULAR without possessing the concept PROPERTY, and thus 6.3.1 is not a priori synthesizable. Nevertheless, he claims, it is a priori because one who possesses all the concepts required to understand the conditional is in position to see its truth.

His point is that one cannot rule out the apriority of a conditional merely on the grounds that one might understand the antecedent but not the consequent. Even if one could plausibly maintain that phenomenal concepts are acquired only via experience, and that the psychophysical conditional thus is not a priori synthesizable, this would not suffice to show that it is a posteriori.

(This example seems poor, as it is doubtful that someone could have the concept of something's being rectangular without having the concept of a property, even if they have not learned the word 'property' or a synonym. Clearer examples of conditionals that are a priori but not a priori synthesizable are easy to find, such as 'if $x = -1$, then $x = e^{\pi i}$ ')

All this is true; however, it is irrelevant to the phenomenal concept strategy. The relevant distinctive feature of phenomenal concepts is not that they are acquired via experience, but rather that they don't bear certain a priori connections to other concepts, and that this explains why someone possessing all the concepts required to understand the psychophysical conditional might not, through reflection on those concepts, be able to deduce its truth.

Assuming we don't go deflationist about reference, phenomenal concept tokens are primitive concepts that denote in virtue of factors that are not cognitively accessible, such as recognitional abilities. As such, their analytic relationship to other sorts of concepts is limited.

Contrast this to the case of 6.3.1. To possess PROPERTY entails, it seems, that one can apply it to instances of the concept RECTANGULAR. Once one possesses the concept PROPERTY in addition to RECTANGULAR, one is in position to see the necessary truth of 6.3.1. Phenomenal concepts, however, don't bear such a priori relations to other concepts. However, this doesn't mean they don't bear any a priori relations.

Stoljar asks us to consider the following:

6.3.2. *If x is a number, then x is not a red sensation.*

Like 6.3.1, 6.3.2 is a priori but not a priori synthesizable. But as Stoljar notes, 6.3.2 is structured similarly to the psychophysical conditional, in that phenomenal concepts are required to understand the consequent and not the antecedent. He concludes from this that the phenomenal concept strategy cannot explain the purported aposteriority of the psychophysical conditional, and so fails.

The phenomenal concept strategy, however, is not an attempt to demonstrate that the psychophysical conditional, if true, is a posteriori. Rather, it is an attempt to explain, through reflection on the nature of phenomenal concepts, how the psychophysical conditional *could be* a posteriori, even if it is necessary. The difference between the psychophysical conditional and 6.3.2 is that there are many a priori arguments to show that 6.3.2 is a priori, arguments that have no analog in the case of the psychophysical conditional. Consider:

(A1) A red sensation exists in time.

(A2) No number exists in time.

(A3) Therefore, if x is a number, then x is not a red sensation.

There's nothing special about this example. One could, for instance, replace 'exists in time' with 'is not divisible by two'. The point is just that the premises of the argument are a priori and the argument is valid. Thus, the apriority of 6.3.2 is demonstrated. There may yet be an a priori argument to show that the psychophysical conditional is a priori as well, but the above argument shows that the apriority of 6.3.2 can be defended without there being an analogous defense for the apriority of the psychophysical conditional.

Note that if Stoljar's argument were sound, it could be applied equally well to the widely accepted 'natural kind concept strategy' of appealing to the special

nature of natural kind concepts to explain how certain statements involving them could be both necessary and a posteriori. It is generally agreed that reflection on the nature of natural kind concepts explains how it is that statements such as ‘If x is a sample of H_2O , then x is a sample of water’ could be both necessary and a posteriori. But consider:

6.3.3. *If x is a number, then x is not a sample of water.*

A natural kind concept is required to understand the consequent of 6.3.3 but not the antecedent, just as it is with ‘If x is a sample of H_2O , then x is a sample of water’. But 6.3.3 is as obviously a priori as 6.3.2, and thus, if we were to follow Stoljar’s reasoning, we would conclude that the natural kind concept strategy fails.

6.4 Stoljar’s Second Argument

Recall that physicalism entails the necessary truth of the psychophysical conditional: ‘If P then P^* ’, where P is a statement summarizing the physical truths of the world and P^* is a statement summarizing the phenomenal truths. Roughly, the conceivability argument against physicalism goes like this:

(B1) It is conceivable that (P and not P^*).

(B2) The conceivability of (P and not P^*) implies the possibility of (P and not P^*).

(B3) Therefore, the psychophysical conditional isn’t necessary.

The physicalist rejects (B2), and employs the phenomenal concept strategy to justify the rejection. But what if this move also rules out other perfectly respectable conceivability arguments? Then it would seem that the phenomenal

concept strategy must fail, on the grounds that it proves too much. This is the crux of Stoljar's second argument.

Stoljar observes that the phenomenal concept strategy posits nothing special about physical concepts and thus would apply equally well to conditionals structured similarly to the psychophysical conditional, but with some other (non-physical) type of concept required for understanding the antecedent. But this would seem to rule out a widely accepted conceivability argument against supervenient behaviorism, i.e., the thesis that phenomenal truths supervene on behavioral truths.

Let B be a conjunction of the behavioral truths of the world and P* be a statement summarizing the phenomenal truths. We may then say that supervenient behaviorism entails the necessity of the *psychobehavioral conditional*: If B then P*. Stoljar asks us to run the conceivability argument against supervenient behaviorism just as we did against physicalism:

- (C1) It is conceivable that (B and not P*).
- (C2) The conceivability of (B and not P*) implies the possibility of (B and not P*).
- (C3) Therefore, the psychobehavioral conditional isn't necessary.

Stoljar notes that if the phenomenal concept strategy refutes (B2) in the conceivability argument against physicalism, then it also refutes (C2) in the conceivability argument against supervenient behaviorism. For if phenomenal concepts are isolated from physical concepts in a way that blocks the move from conceivability to possibility, then they are similarly isolated from behavioral concepts.

It is true that if the phenomenal concept strategy blocks (B2) in the conceivability argument against physicalism, then it also blocks (C2) in a similarly structured argument against behaviorism. But the conceivability argument against

supervenient behaviorism needn't be structured in this way, nor need it apply to all forms of supervenient behaviorism. Very roughly speaking, we can, for present purposes, distinguish two strains of behaviorism that accept the psychobehavioral conditional: call them a priori behaviorism and a posteriori behaviorism. Let us deal with each in turn.

(We need only concern ourselves with those behaviorist theories that endorse the psychobehavioral conditional. Thus, we needn't worry about eliminative behaviorism, which holds that there are no phenomenal truths, or methodological behaviorism, which makes no metaphysical claims.)

A priori behaviorism, sometimes known as logical, analytic, or philosophical behaviorism, is the thesis that mentalistic terms or concepts (including phenomenal terms or concepts) can be analytically reduced to behavioral terms or concepts. When we attribute a phenomenal state to someone, we are, according to a priori behaviorism, characterizing the person a priori as being in some state of overt behavior or as having certain dispositions to overt behavior. The theory is most prominent in the work of Carnap and Ryle. I take the following to be a sound conceivability argument against a priori behaviorism:

(D1) It is coherently conceivable that (B and not P*).

(D2) Therefore, the psychobehavioral conditional is not a priori.

Unlike the conceivability argument against physicalism, it is not a premise of this argument that conceivability entails possibility.

Now let us turn to a posteriori behaviorism, which holds that a subject's phenomenal states are a posteriori identical to or supervenient on complex dispositional states that can, in theory, be adequately characterized in terms of the subject's overt behavior and external physical environment. There is no obviously sound conceivability argument against a posteriori behaviorism, just as

there is no obviously sound conceivability argument against a posteriori physicalism. Rather, I take it that a posteriori behaviorism should be rejected on more mundane grounds, namely that it is explanatorily inadequate, especially in contrast to a posteriori physicalism.⁶

Although Phenomenal Empiricism is false (knowing what it is like to have a phenomenal property does not require that one have had the property), the phenomenal concept strategy is none the poorer. We have been given no reason to abandon the appeal to special features of phenomenal concepts to explain why the leading dualist arguments fail.

⁶There have been many explanatory inadequacy attacks on a posteriori behaviorism. The *locus classicus* is Chomsky (1959).

Chapter 7

The Property Dualism Argument: Reply to White

7.1 Introduction to the Argument

The so-called property dualism argument has received considerably less press than the related conceivability and knowledge arguments. The reason for this is most likely that while the conceivability and knowledge arguments have as their targets all flavors of physicalism, including many forms of functionalism, the property dualism argument is aimed only at type-identity theories: i.e., those theories that take phenomenal properties to be identical to physical properties (and which hold such identities to be a posteriori).

The property dualism argument is worth paying special attention to, however, because against identity theories the argument has a key virtue over the conceivability and knowledge arguments. The virtue is that the argument in no way requires metaphysical possibility to follow from conceivability. Thus it sidesteps (for the most part) long-fought battles in the areas of modal epistemology and the metaphysics of modality. Perhaps because of this, the argument has generated renewed interest and vigorous debate.¹

The classic property dualism argument is, I think, convincingly answered by

¹Its most careful contemporary articulation and defense can be found in (White, 2006), and notable responses to the argument can be found in (Perry, 2001) and (Block, 2006).

the phenomenal concepts strategy. But recent refinements to the argument, especially by White (1986, 2006), overcome a simple appeal to phenomenal concepts. Some physicalists, such as Block (2006), have attempted to respond to the refined version, but such responses rely crucially on a quotational account of phenomenal concepts which we found inadequate (see chapter two).

7.1.1 The Classic Property Dualism Argument

Contemporary discussion of the property dualism argument traces to Smart's (1959) defense of the identity theory. Among the objections to the theory is one credited to Max Black:

Suppose we identify the Morning Star with the Evening Star. Then there must be some properties which logically imply that of being the Morning Star, and quite distinct properties which entail that of being the Evening Star.²

We might unpack Black's objection as follows: In order for an identity to be a posteriori, it must be possible to rationally doubt its truth. But this can happen only if the terms flanking the identity pick out their referents in different ways (Frege's constraint), so that one cannot simply derive sentences employing one of the terms from sentences employing the other in its place, nor from sentences employing the other term plus sentences expressing one's background knowledge (assuming the background knowledge does not itself entail the identity).

If, for example, a putative identity such as *Pain = C – fiber stimulation* is true a posteriori, it must be the case that the way we pick out pain is different from the way we pick out C-fiber stimulation. These different ways must be such to underwrite the contention that the presence of pain is not derivable a priori from the presence of C-fiber stimulation and background knowledge.

If we understand “ways of picking out” an object as latching on to the object via two or more distinct properties of the object, then the physicalist faces

²(Smart, 1959), p. 148

a dilemma. For if ‘Pain’ picks out its referent via its quality of painfulness and ‘C-fiber stimulation’ picks out its referent via its quality of being a C-fiber stimulation, then painfulness and the quality of being a C-fiber stimulation will either have to be distinct properties (in which case we are saddled with property dualism), or there will have to be some further distinct properties to explain the aposteriority of the property identity *Painfulness = being a C – fiber stimulation*. Obviously, this is the start of a vicious regress. In either case, we are left, on pain of incoherence, with some sort of property dualism or dual aspect theory (even as we grant the truth of *Pain = C – fiber stimulation*).

In response to Black’s objection, Smart asserts that introspective phenomenal reports are *topic-neutral*. Smart writes “When a person says ‘I see a yellowish-orange after-image’ he is saying something like this: ‘*There is something going on which is like what is going on when I have my eyes open, am awake, and there is an orange illuminated in good light in front of me*’”.³

Smart’s response is inadequate, for at least the following reasons:

“What is going on” does not uniquely identify an experience type. There are *indefinitely many* things going on when I have my eyes open, etc. The topic-neutral approach does not explain how we can introspectively pin down a type of subjective feel and formulate thoughts about it as opposed to other property instances of which the subject is simultaneously aware.

We can type experiences by their subjective feel or phenomenal character, independently of the conditions surrounding their instantiations. We may, for instance, be wrong about or unaware of what (non-phenomenal) events are occurring simultaneous with the phenomenal property instance. However, this does not seem to in any way undermine our ability to pick out the phenomenal type in question.

³(Smart, 1959), p. 149

We conceive of phenomenal properties as genuine properties which we can think about, ponder apparently substantive similarities and differences, and embed in complex thoughts independently of the conditions surrounding their instantiations. For example, I can think about what it's like to have an experience of red without thinking at all about the conditions surrounding past experiences of the kind in question.

For these reasons, and others, a better reply to Black's objection is required.

7.1.2 The Phenomenal Concept Reply

The phenomenal concept strategy provides for a response to Black's argument that is much like that provided in response to the knowledge and conceivability arguments. We might employ the now-familiar strategy of asserting that what appears to be a metaphysical distinction is merely a conceptual distinction.

We can accept Frege's constraint, which holds that the terms flanking the identity must pick out their referent in different ways, but deny that these "different ways" involve latching on to distinct properties of the object. Rather, we can assert that the "different ways" are to be expressed as different ways of conceptualizing the object of the identity statement at the personal level. In other words, the difference lies in the types of concepts or representations in the subject's mind, rather than in distinct properties of the object of the identity statement itself.

Thus there is a ready explanation of how one can rationally believe a pain both to exist (through a phenomenal concept) and not exist (through a physical concept) or vice versa — the physical and phenomenal concepts co-refer, but this co-reference is not deducible a priori. Thus we have the now-familiar "two concepts, one state" response.

Rephrasing this in terms common to recent literature on the property dualism

argument, let us call these concepts “representational modes of presentation”, and the distinctive properties of the objects that the representational modes of presentation represent, “metaphysical modes of presentation”.

What we call “metaphysical modes of presentation”, following Block, White calls “non-representational modes of presentation”. The relationship between these phrases and Frege’s original use of “mode of presentation” is tricky. I will be concerned only with contemporary usage.

I will assume that representational modes of presentation are identical to concepts. Some may balk, on the grounds that representational modes of presentation have functional roles that are distinct from those of concepts. This is not crucial, so long as each representational mode of presentation is essentially associated with a specific concept type. To my knowledge, none of the literature claims that anything rides on such a distinction.

The contemporary a posteriori physicalist holds that in the case of *Pain = C – fiber stimulation*, there are two representational modes of presentation but only one metaphysical mode of presentation, and that it is the distinctness of the representational modes of presentation and lack of relevant a priori connections between them that explain the a posteriori character of the identity. Since there is only one metaphysical mode of presentation, physicalism seems not to be threatened.

7.2 The New Property Dualism Argument

I now turn to examine an influential contemporary refinement of the property dualism argument by White (2006). White makes many refinements to, and elaborations of, the argument, but the key point that will concern us here is the proper explication of the seemingly analytic claim that if a statement is true a

posteriori, it must be the case that a subject could rationally doubt its truth. For White, this amounts to coherently conceiving a situation in which the statement is false. He argues that unless at least one side of the identity is picked out via a contingent mode of presentation, this cannot be done.

The identity theorist must deny this, on pain that any true psychophysical identity be a priori. The challenge for the a posteriori physicalist is to make the case that the nature of phenomenal concepts explains how a perfectly rational subject might doubt a true psychophysical identity *without* being able to coherently conceive of a situation in which the identity is false. This is the challenge I take up here.

7.2.1 Representational Coherence

The worry pressed by White is that the phenomenal concepts reply cannot save the identity theorist from incoherence. Because his physicalism is a posteriori, he must offer a story of how doubt in the psychophysical conditional is possible. But because he holds the psychophysical identities to hold necessarily (on the standard Kripkean theory), he cannot explain the doubt by way of (merely) epistemically possible worlds in which the identity statement is false. For the physicalist cannot coherently describe such a world.

To see the problem in clearer focus, we might ask what the physicalist who rationally doubts the truth of an identity such as *Pain = C – fiber stimulation* would take a world in which it is false to be like. If a sentence S is true a posteriori, one might reasonably suppose that there is some fully describable coherent situation in which -S is true.⁴ After all, if there were no such coherent

⁴I use 'situation' as opposed to 'world' or 'possible world' to avoid any implication of metaphysical possibility. Also, 'fully describable' should be understood in a loose way which does not go beyond the uncontroversial ways we take metaphysically possible worlds to be describable.

situation, then it seems S would be true a priori. But just how are we to coherently describe a situation in which there is, say, C-fiber stimulation but no pain, given the identity theorist's contention that C-fiber stimulation and pain refer via the same metaphysical mode of presentation?

White's point is that conceivability of the negation of an identity requires more than just distinct representational modes of presentation for the terms flanking it. What is needed is a requirement more robust than Frege's, one which ties conceivability to a full and positive conception of how things might be. White calls this the requirement of *representational coherence*:

Representational Coherence. *the requirement that there be a coherent account of what the world would be like if it were the way that such an uninformed or misinformed (but still fully rational) subject took it to be.*⁵

Obviously, the existence of distinct representational modes of presentation is not sufficient for establishing an identity to be a posteriori. Consider, for example, a complex true mathematical identity statement. There is an clear sense under which there are distinct representational modes of presentation for the terms flanking the equality operator, and a reasonable but uninformed person could doubt the truth of the identity. Still, the identity is a priori, and one cannot rationally conceive its negation.

The appeal to the lack of a priori connections between representational modes of presentation is not enough to account for the a posteriori character of psychophysical identities, since representational modes of presentation, by themselves, do not specify content. We have to look at the associated metaphysical modes of presentation to determine the content of a belief.

In standard cases of a posteriori identity, there are two metaphysical modes

⁵(White, 2006), p. 212

of presentation which correspond to the identified terms. Take, for example, the identity *Morning Star = Evening Star*. In this case, one can rationally doubt the truth of the identity by considering a situation in which the object that bears the property of being the morning star does not bear the property of being the evening star. This route, however, is unavailable to the a posteriori physicalist who must deny that Pain and C-fiber stimulation refer via distinct properties of the referent.

Could one just reject representational coherence? It seems not. Beliefs have *content*, and there needs to be a story of how a belief has the content it has — a story which must go beyond that belief’s functional role within an internally specified system. Such a story will inevitably draw a connection between factors internal to the subject and properties external to him. To deny such a story is, as White notes, to be guilty of a vehicle-content confusion — to relegate the content of beliefs (and of thoughts generally) to the status of an uninterpreted calculus. To deny representational coherence in defense of the phenomenal concept strategy is to bite far too big a bullet.

7.2.2 The “A Priori Condition”

What seems to explain the subject’s ability to rationally doubt *Morning Star = Evening Star* is, minimally, his ability to positively conceive of a scenario, consistent with his evidence, in which the sentence is false. What makes this possible, given the truth of the sentence, is that there are distinct representational modes of presentation associated with ‘Morning Star’ and ‘Evening Star’ which are tied a priori to distinct metaphysical modes of presentation (i.e., the properties of being the Morning Star and being the Evening Star, respectively).

Thus the subject is free to positively conceive a scenario in which the morning star is not the evening star. Even if such a scenario is a skeptical one, given the

subject's evidence, it is not a priori so long as it does not *contradict* the subject's evidence (and as any first year philosophy student learns, there are all sorts of outlandish scenarios that are compatible with one's evidence).

If this story of how one can come to doubt a true identity is true for *all* a posteriori identities, then the following thesis holds:

A Priori Condition. *In order for a belief statement to represent things as being a certain way, there must be an a priori connection between the terms of the statement, the associated representational modes of presentation, and the properties by which those terms picks out the states they do.*

The Condition has its fans on both sides of the physicalism debate, but it is quite dangerous for the physicalist to accept it. For if it is true, then the phenomenal concept strategy seems to fail to adequately explain how one can rationally doubt psychophysical *property* identities, since the avenue of claiming one object but two properties, as in the morning star-evening star case, is closed. One can have a robust account of phenomenal concepts under which phenomenal concepts acquire their reference solely in virtue of sub-personal causal facts.

One notable physicalist who supports the A Priori Condition is Block (2006), who rebuts the property dualism argument by challenging the supposition that the Condition implies that metaphysical modes of presentation must be "thin", i.e. expressible by predicates whose intensions do not depend on empirical facts about the actual world. If representational modes of presentation can be associated a priori with metaphysical modes of presentation that are not thin, then there is room to conceive of alternative worlds in which the metaphysical mode of presentation has different essences. Block's case, while interesting, relies crucially on his quotational theory of phenomenal concepts, and is thus of no use to a physicalist who finds the quotational account unacceptable (see chapter three of this dissertation).

It seems to me that the physicalist's best strategy is to accept the requirement of representational coherence but deny the a priori condition: that is, to accept that an identity statement is a posteriori only if some subject can rationally doubt its truth, to also accept that in order to rationally doubt an identity statement one must be able to coherently represent its negation as being true, but to deny that all the properties of the represented scenario are knowable a priori.

Rather, the physicalist should claim that the aposteriority derives from the subject's realization that reference-fixing for phenomenal concepts is at least partly a matter of subpersonal causal mechanisms, and that the represented scenarios include indexical placeholders for properties partly determined by such mechanisms. For any subject and any phenomenal concept, there are scenarios consistent with one's evidence in which such mechanisms might have picked out a property different from the physical property picked out in the actual world. (To be clear, this is not to make the claim that phenomenal concepts fail to denote rigidly).

Thus, the content of the represented scenario in which the psychophysical identity in question is false is given not by a positive account of the relevant properties instantiated in the situation in question. Rather the content is given, in part, by semantic placeholders (like indexicals) whose referents are not specified by the semantic placeholders themselves, nor by any other term or sentence in the scenario description. The placeholders get their reference from factors which are not part of the scenario description, but which are fully determined by certain subpersonal facts about the subject (and perhaps his environment).

White would not be happy with this defense, as he argues that subpersonal factors have no bearing on a subject's rationality. If his claim is true, then it is irrelevant whether a subject could, in this way, successfully specify a scenario in which the psychophysical identity is false. For such a specification, even if

successful, would do nothing to explain the subject's rational doubt in the truth of the identity.

To drive home this claim, he asks us to consider the following fanciful case:

Imagine on the subject of Jones's honesty, Smith is genuinely irrational. Smith believes in some contexts that Jones is fundamentally honest and trustworthy and in others that Jones is fundamentally dishonest and not to be trusted [...] Suppose that, without Smith knowing it, he is in contact with two distinct people Jones and his dishonest twin. And suppose that in those contexts in which Smith is disposed to believe that Jones is honest, Smith is almost always in contact with Jones, and that in those in which he is disposed to believe the contrary, he is in causal contact with Jones's dishonest double. Would this in any way undercut the assessment of Smith as irrational on the subject of Jones's dishonesty? The answer is no. [...] Thus it seems clear that we cannot rationalize an otherwise irrational belief set by appealing to subpersonal functional states and external causal chains to which the subject has no access.⁶

We can, I think, grant that if phenomenal concept possession involves nothing over and above the possession of certain recognitional abilities, then White's point is decisive. For suppose, by analogy, that Smith's concept of Jones were purely recognitional in this sense. On such a view, Smith would, by virtue of external factors, possess distinct recognitional concepts one for Jones and one for Twin-Jones. Smith's thoughts would thus be consistent in that those which hold Jones to be honest would not contradict those which hold Twin-Jones to be dishonest. But it is obvious that Smith's thoughts are irrational.

Against White, one might claim that the problem in the Smith case is that it relies on luck. Smith does not really have the ability to recognize Jones nor the ability to recognize Twin-Jones. Recognitional abilities, one might claim, must pass some sort of modal or counterfactual test. As presented, it seems obvious Smith would fail such tests, whatever they may be, as his success at re-identifying Jones and Twin-Jones is a matter of dumb luck. Thus one might reasonably claim

⁶(White, 2006), p. 218

that Smith lacks recognitional concepts of both Jones and Twin-Jones, and so is irrelevant to the issue at hand.

Such a response, however, would seem to miss the point. For we can change White's story slightly to make Smith's ability to discriminate Jones from Twin-Jones more modally robust, without in any way mitigating Smith's irrationality. For suppose that Jones, but not Twin-Jones, possessed some property (a slight odor, say) that is causally responsible for Smith's feelings of distrust but is itself completely unknown to Smith. In such a case, Smith would reliably and robustly discriminate Jones from Twin-Jones, but not in a way that mitigates the irrationality of trusting one but not the other.

7.2.3 Recognitional Concepts vs. (Mere) Recognitional Abilities

I think the physicalist's best response to the sort of worry White raises, i.e., the worry that motivates the a priori condition, is to explicate phenomenal concept possession conditions in a way which involves more than recognitional abilities. I suggest that White is overemphasizing the role recognitional abilities play in the individuation of phenomenal concepts.

While it is the view defended in this work that phenomenal concepts are essentially recognitional concepts, this should in no way to be understood as endorsing the claim that concepts are abilities, either in general or with respect to phenomenal concepts alone. I hold that concepts are mental representations (or essentially involve mental representations). What makes recognitional concepts special is that their referents are determined by factors extrinsic to the representation, and typically by factors extrinsic to the thoughts in which such representations play a role. While it is true that recognitional abilities count

among the (essential) possession conditions of phenomenal concepts, they by no means exhaust such conditions.

As a realist about mental representation, I hold that when a concept is deployed, there is a real instance of the representation. That is, there is an event or state that is an instance of it. As a physicalist, I hold that these events or states are physical events or states. Thus, a basic possession condition for a phenomenal concept, as for any concept, is that a subject must be able to bring about particular sorts of (presumably brain) events. Such a requirement is over and above the requirement that the subject have the ability to introspectively recognize some phenomenal property.

What individuates phenomenal concepts are not merely their referents, which are determined subpersonally, but also the nature of the representations themselves. I have argued that phenomenal concepts are representationally simple, in the sense that they are not constituted by simpler concepts. But to say that representations are simple is not to say that they are intrinsically indistinguishable (though perhaps this is the received Lockean view).

Rather, it is perfectly consistent to suppose that although these representations do not decompose to simpler representations, they nevertheless have some intrinsic features that serve to identify their instances as being instances of some particular type of representation. Nor is it to deny that such intrinsic identifying features are available to the subject at the personal level.

To illustrate, let us consider two subjects, Brown and Johnson, distinct phenomenal concepts, C_1 and C_2 , and their referents Q_1 and Q_2 . Suppose that Brown, in the past, possessed both C_1 and C_2 , but now unknowingly has lost the ability to distinguish Q_1 from Q_2 (and thus unknowingly has lost possession of both C_1 and C_2). Brown will still distinguish two concept types at the personal level (i.e., Brown still thinks he can think separate Q_1 and Q_2 in thought), even though he

now possesses neither. If such a subject came to contradictory judgments about the union of Q_1 and Q_2 , we would not hold him to be irrational. Conversely, we could imagine that Johnson possesses a phenomenal concept, C_{12} that he deploys reliably in response to the union of Q_1 and Q_2 , and who possesses no other phenomenal concepts of Q_1 or of Q_2 . Unlike Brown, however, Johnson could not, on pain of irrationality, come to the same contradictory judgments as Brown.

In the story above, Brown and Johnson have identical recognitional abilities with respect to the union of Q_1 and Q_2 . But they do not possess the same phenomenal concepts. Johnson possesses a single phenomenal concept (of the union of Q_1 and Q_2), and Brown possesses two concepts, neither of which is properly a phenomenal concept (as neither is deployed reliably in response to any phenomenal property).

We are now in position to see that White's case of Smith and Jones has no teeth against a representationally realist theory of phenomenal concepts. For the proponent of such a theory can agree with White that questions of the rationality of judgments involving phenomenal concepts must turn on factors accessible at the personal level, without denying that the reference of phenomenal concepts is determined by factors not accessible at that level.

Those who are deeply skeptical of the existence of mental representations are often drawn to the idea that concepts can be explicated entirely in terms of abilities. If all a concept of Q amounts to is, say, the ability to attend to Q , then one can formulate a theory of concepts without mental representations. And a theory of phenomenal concepts, on this approach, would just amount to a theory of some specialized abilities, such as the ability to attend to a phenomenal property in introspection. It seems that White had this sort of position in mind as the target of the Smith and Jones example.

On the one hand, we have robust concepts with all sorts of a priori connections

to the properties they denote (e.g., functional concepts). On the other hand, we have mere recognitional *abilities*, which in my view are not concepts at all. Phenomenal concepts occupy a middle ground, as mental representations with a simple structure.

I think that general considerations of conceptual productivity and systematicity are decisive against the identification of concepts with abilities (or, equivalently, concept possession with ability possession). However, one can still accept a representational theory of mind that includes some recognitional concepts, i.e., concepts whose possession conditions essentially include some recognitional abilities (see chapter four of this dissertation).

It seems that White presents the physicalist with a trilemma:

- (1) Deny representational coherence. Problem: To deny representational coherence is to be guilty of a vehicle-content confusion. It's unclear that, without representational coherence, beliefs could have content at all.
- (2) Accept the a priori condition. Problem: The a priori condition is incompatible with a posteriori physicalism, as one cannot positively coherently conceive a scenario in which a true psychophysical identity is false.
- (3) Deny the a priori condition and individuate phenomenal concept types solely via subpersonally determined recognitional abilities. Problem: As White's case of Smith and Jones shows, one cannot account for the rational doubt aposteriority requires by appealing to subpersonal factors.

I suggest this trilemma is a false one. Instead I suggest:

- (4) Deny the a priori condition and individuate concept types at the personal level, via factors intrinsic to their instances. Hold that although the reference of a phenomenal concept is determined by sub-personal factors, such

factors are irrelevant to determinations of the rationality of a judgment. Rather, such determinations will depend on factors available to the subject, factors intrinsic to the representations being tokened. Hold that tokens of different types of simple mental representations can be distinguished by the subject regardless of the properties (if any) they represent.

It is true, as White asserts, that appealing to external causal connections cannot rationalize a subject's beliefs or doubts. This is true, but unimportant. Content is determined by causal connections, but it is the individuations and interconnections at the personal level that attest to rationality or lack thereof.

7.3 The Importance of Type-Identity to A Posteriori Physicalism

The phenomenal concept strategy is not limited to a defense of type-identity physicalism. Although this dissertation has been written with type-identity in mind and examples have been chosen which assume that such a position is at issue, it is not a central component of the phenomenal concept strategy. The strategy requires only a certain realism about phenomenality, such that phenomenal concepts can successfully denote phenomenal types. It is mum, however, on the nature of these properties.

The phenomenal concept strategy could be of use, for example, to a psychofunctionalist or even an a posteriori dualist. That is, the main effect of the strategy, if successful, is to explain and defend the lack of a priori connections between the physical and the phenomenal realms. Just what the referents of the phenomenal concepts actually *are* is not something to be decided by the strategy. Rather, other considerations, such as theoretical simplicity, elegance, and

productivity, must come into play.⁷

So in order to motivate concern for the property dualism argument, an argument which has been relatively neglected next to the conceivability and knowledge arguments, I'd like to briefly discuss why a defense of type-identity is important for the viability of an a posteriori physicalism and thus why a physicalist should care about the property dualism argument.

Up until the mid 70's or so, to be a physicalist was, minimally, to hold that mental types (e.g. pain) are identical to and re-describable as physical types (e.g. C-fiber stimulation). Following convention, I will refer to the family of such theories as 'the identity theory' .

The identity theory is distinct from the thesis of the reducibility of psychology. While the former asserts that mental properties can be re-described as physical properties , the latter asserts (at a minimum) that the truths of psychology can be re-described as truths of the hard sciences .

For present purposes, I will assume that the anti-reductionist rejects both the identity theory and the thesis of the reducibility of psychology. Since any problem with the identity theory is also a problem for the thesis of the reducibility of psychology (due to its need for vocabulary-connecting bridge laws), but not vice versa , I will from this point characterize the anti-reductionist simply as one who rejects the identity theory.

Proponents of the identity theory saw it as a compelling move toward a naturalistic integration of the mental into the scientific domain. It supposed neither dualism's queer entities and forces, nor behaviorism's peculiar methodological and ontological strictures.

But perhaps the nicest thing about the identity theory is that it provided a physical account of mental causation. For if mental properties are re-describable

⁷This point is made more elaborately in chapter five.

as physical properties, and physical properties are causally relevant, then mental properties are causally relevant.

Despite these virtues, the identity theory has serious drawbacks. First, one may question the theory's optimism. When one traverses the history of science, one finds precious few clear instances of reduction. Even celebrated cases, such as the putative reduction of the classical gene to molecular biology, have come under intense scrutiny. Perhaps the belief that mental types can be precisely identified with physical types is an unwarranted pipe dream.

Also, there are strong intuitions that mental states are multiply realizable. If it is possible that some species without c-fibers can feel pain, then pain cannot be identical to C-fiber stimulation. If these intuitions are correct, then mental types are at best identical to ad hoc infinite (or at least exceedingly large) disjunctions of physical properties.

The identity theory can also be too liberal in attributing mental states. To see this, consider that if pain is C-fiber stimulation, then it follows that we can create pain simply by stimulating excised c-fibers in a Petri dish.

Physicalists were drawn to the identity theory because it provided an ontologically modest account of mental states that (contra analytic functionalism) made them concrete, causally relevant, and thus a part of the ordinary domain of scientific inquiry. So when token identity theories popped onto the scene, claiming all of the identity theory's virtues and none of its vices, they quickly became the hottest ticket in town.

The token identity theory holds mental particulars to be identical to physical particulars. It does not require that mental types be identical to physical types. The theory rose to prominence on the coattails of Donald Davidson, whose theory of anomalous monism gained scores of converts.

At first glance, token identity seems to do it all. It skirts the thorny issues

of reduction, as it does not require that psychophysical reductions exist. It has no problem with multiple realizability, as it claims no physical similarities among tokens of a mental type. And it seems to provide a straightforward account of genuinely causal psychophysical interaction (we will return to question this). For if mental particulars are identical to physical particulars, and physical particulars can be (singular) causes, then mental particulars can be (singular) causes. An experience of pain has causal powers over the physical because it is identical to a firing of c-fibers, which has causal powers over the physical.

Taken by itself, token identity isn't much of a physicalist theory. For it is compatible with the possibility that two physically indistinguishable particulars, P1 and P2, are respectively identical to distinguishable mental particulars, M1 and M2. Token identity allows that two physically identical brains, side by side, could differ mentally in any way you'd like. This hardly seems to vindicate physicalism! The physicalist needs to bolster token identity with a stipulation that rules out such possibilities. Enter supervenience.

At a minimum, the physicalist should be inclined to supplement token identity with what Kim (1984) has called weak supervenience:

Weak Supervenience. *A weakly supervenes on B if and only if necessarily for any x and y if x and y share all properties in B then x and y share all properties in A - that is, indiscernibility with respect to B entails indiscernibility with respect to A.*

Weak supervenience is compatible with anti-reductionism. For it doesn't require that the physical properties of objects fix their mental properties. All it requires is that indiscernible physical objects in the world be indiscernible mentally.

Under weak supervenience, c-fiber firings may be pains in this world, but be pleasures or nothing mental at all in physically indiscernible worlds. Since laws,

on standard accounts, are supposed to say things about at least some worlds other than our own (they underwrite true counterfactual conditionals, for example), weak supervenience, which legislates restrictions within a single world only, doesn't imply any law-like psychophysical regularities.

Fortifying supervenience requires adding another modal operator. Kim defines strong supervenience thusly:

Strong Supervenience. *A strongly supervenes on B just in case, necessarily, for each x and each property F in A, if x has F, then there is a property G in B such that x has G, and necessarily if any y has G, it has F.*

Strong supervenience takes us closer to standard physicalist intuitions, but it may not be compatible with anti-reductionism. Kim (1984) that strong supervenience implies that every supervenient property in A has a necessary coextension in B. Under strong supervenience, then, there are necessary property to property entailments between the mental and the physical. Are these entailments properly called laws? Do they constitute a reduction? I'm not sure, but I think so.

The entailments are like laws in that they express necessary connections and are counterfactual-supporting. They are unlike what we ordinarily call laws in that they may not be expressible in a language and are unlikely to be available for epistemic activities such as explanation and reduction.

These disanalogies are not critical. Our concern here is with the way the world is - not with our ability to formulate its truths or use them for our purposes. Let us put aside these tangential concerns. If Kim's analysis is correct, then we have no good reason to deny that strong supervenience implies the existence of psychophysical laws.

Hence strong supervenience is off-limits to the anti-reductionist, even though it doesn't entail the existence of the vocabulary-bridging laws prized by the identity

theory.

What is it about particulars A and B that make it true that A causes B? A prominent account, endorsed by Davidson among others, explains causation in terms of laws. Here's one formulation:

A causes B just in case A is a token of type A*, B is a token of type B*, and A and B fall under a strict law connecting A* and B*.

Suppose A is a mental particular and B is a physical particular. If A is the cause of B, then, by the nomological account, A is the cause of B in virtue of some property A*, that connects to some property B* via a strict law.

Given the causal closure of the physical, we have a physical property A* that, along with background conditions, implicates B* via a strict law. Thus we have it that A causes B in virtue of A's physical properties. All the causal work is done. What work left is there for A's mental properties to do? Kim calls this the exclusion problem.

Davidson (1993) has denied the exclusion problem. For although he is a proponent of the nomological account of causation, he maintains what he calls a purely extensionalist view of the causal relation. Causation, for him, is a two-place relation between events - however described. That physical events have physical descriptions is irrelevant to their causal efficacy. To Davidson, it makes no literal sense to say speak of something being a cause in virtue of one of its properties.

Davidson's response to exclusion problem is lacking. As McLaughlin (1993) has pointed out, one can consistently hold both i) that the relata of the causal relation are concrete particulars (or that in singular causal statements, 'cause' is a two-place predicate subject to substitutivity of identicals and existential generalization), and that ii) that if A caused B, it did so in virtue of some of A's properties.

Take, for example, the relation ‘weighs less than’. This is a two-place predicate that admits of substitutivity of identicals. It still is the case that if A weighs less than B, it does so in virtue of some property - namely weight.

Davidson could claim that causation is a brute relation between events. According to this view, there simply are causal facts, and these facts are not in any sense logical constructs from other (non-causal) facts, though they can be inferred from them. This seems immensely unappealing as it abandons the ontological spirit of physicalism by leaving the world haphazardly populated by causes that are undetermined by the world’s physical state and history.

Perhaps one should weaken the nomological account of causation to allow for non-strict causal laws. On this modified account, generalizations that require *ceteris paribus* clauses could count as causal laws. The advantages here are obvious, as mental events could now be causally implicated in virtue of their falling under mental types. This could be done without retreating to strong supervenience, and so reductionism could be avoided.

The costs of this approach are unclear, as much remains to be said about the nature of non-strict laws - particularly in spelling out the *ceteris paribus* clauses. If *ceteris paribus* clauses are in some deep sense context-dependent or are only pragmatically determinable, then the extensionality of causation is threatened. If, however, there are no such implications, then non-strict laws may be the most promising route.

Even if non-strict laws could be made respectable, we might still be faced with a dilemma. Given that the causal relation between two events is grounded in the basic physical properties of the events and a strict law relating them, what causal work is there left to do? It seems that either the non-strict law reduces to a strict law in the particular case (on this view, non-strict laws are merely strict laws with some of their antecedent conditions existentially quantified), or it does

not.

If the non-strict laws do reduce, then concrete events are really just related by strict laws, and we are back to square one. If the non-strict laws fail to reduce, then we are stuck with competing (strict) physical and (non-strict) psycho-physical causal accounts, and it is difficult to see what the psycho-physical account contributes to the account other than an epistemic aide for our limited minds.

Another way to try to get at mental causation under non-reductive weak supervenience is to promote an alternative notion of causal relevance. Loewer and LePore distinguish two forms:

Properties F and G are causally relevant¹ to c's causing e just in case c's having F and e's having G makes it the case that c causes e.

They are relevant² when c's being F brings it about that e is G.

They argue that anomalous monism (and, I suppose, non-reductive weakly supervenient theories generally) allows mental properties relevance², even if it denies them relevance¹. If one finds relevance² to be a compelling property, perhaps the charge of epiphenomenalism is mitigated. McLaughlin, however, doubts that many have conflated the two notions of relevance. Rather, he thinks that they are only interested in relevance¹.

Accepting the nomological account of causation puts the anti-reductionist physicalist in a tight spot. Token identity plus weak supervenience allows one to escape the specter of reductionism, but at the steep price of prima facie causal irrelevance for mental properties. On the other hand, strong supervenience, as Kim has shown, is reductionist in that it implies the existence of necessary entailments between mental and physical properties.

This may not be so bad, as strong supervenience does not imply the existence of logical entailments between mental and physical *terms*, nor does it imply that any of the implied entailments are knowable or expressible. If one felt pushed to

anti-reductionism, not by a priori arguments, but by skepticism over the existence of vocabulary-connecting bridge laws and by other doubts about the possibility of theoretical reduction, then perhaps one shouldn't fear strong supervenience and its mild reductionist implications. The remaining anti-reductionists should probably look for an alternative to the nomological account of causation.

7.4 An Additional Problem for Type-Demonstrative Accounts of Phenomenal Concepts

Those who take phenomenal concepts to refer via a demonstrative introspective act (at least in canonical cases) sometimes claim that phenomenal concepts refer *directly*, and thus that there is no mode of presentation on the phenomenal side, only sub-personal causal links that serve to establish reference.

Sometimes, a similar point is made by claiming that phenomenal states serve as their own modes of presentation. This is problematic, however, insofar as we are taking modes of presentation to refer to properties of the referent, rather than the referent itself (i.e. a type of experience). If this is true, then the physicalist has a way to describe a situation which captures the thought that C-fiber stimulation is not identical to pain, i.e. a situation in which subpersonal causal connections determine the referent of 'pain' to be something other C-fiber stimulation.

This will not work, however, because demonstrative reference does require an metaphysical mode of presentation. Furthermore, there is a strong case to be made that demonstrative reference frequently, if not always, involves an representational mode of presentation that is associated to the metaphysical mode of presentation a priori.

To see this, consider two demonstrations made of a single object - call the referents D1 and D2. Now, there are many cases in which it is perfectly rational

to doubt the identity of D1 and D2, namely those cases where the two demonstrations refer to an object via different properties of that object. (Such as first demonstrating an airplane via the sight of its nose, and then via the sight of its tail).

What rationalizes the doubt is that the subject is aware that she has picked out an object via distinct representational modes of presentation, and that these representational modes of presentation are associated a priori with different metaphysical modes of presentation. She can thus easily imagine the associated metaphysical modes of presentation belonging to different objects (Such as the demonstrated airplane nose belonging to a different plane than the demonstrated tail). The subject's doubt is thus rationalized in fundamentally the same way as it was in the case of the morning star and evening star - via distinct properties of the object demonstrated.

What separates demonstrative reference from descriptive reference is not an absence of representational modes of presentation in the former, but rather that the representational mode of presentation of a demonstrative is not descriptive in nature. As far as the present argument goes, this is a distinction without a difference. Doubts of the truth of demonstrative identities are generally rationalized by distinct representational modes of presentation which are tied a priori to distinct metaphysical modes of presentation. This will not do for the physicalist.

7.4.1 Phenomenal Concepts, Demonstratives, and Pure Indexicals

In this section, I suggest that non-demonstrative indexicals do not face this problem, as they do refer without an representational mode of presentation. If this is true, then that is another point in favor of the account of phenomenal concepts

defended in chapter three. After all, no concept that violates Phenomenal Empiricism could be a demonstrative, as it is essential to a demonstrative that its content is what is being (or has been) demonstrated. And this, of course, requires the existence of what the concept denotes (or an instance of the denotation, in the case of a type-demonstrative).

Kaplan (1989) distinguishes between what he calls “true demonstratives” and “pure indexicals”. The reference of an utterance of a true demonstrative is determined (in part) by the speaker’s accompanying actions or intentions. For example, the reference of an utterance of ‘that’ is determined (in part) by the speaker’s accompanying pointing gestures, or by the speaker’s intention to refer to a particular object.

The reference of a pure indexical, in contrast, is not determined by the speaker’s actions or intentions in this way. For instance, an utterance of ‘I’ refers to the speaker, whether or not she points at herself, and an utterance of ‘tomorrow’ refers to the day after the day of utterance, regardless of the speaker’s intention to refer to some particular day. We can say (loosely speaking) that the reference of pure indexicals is *automatic*, whereas the reference of true demonstratives requires something extra from the speaker.

Consider, for example, the indexical ‘here’. I am, I believe, in New Brunswick, NJ. So I believe that *here = New Brunswick*. But this is not through any mode of presentation. I could, for instance, be in a sensory deprivation tank and as easily form the same belief. I might find it difficult, however, to token a demonstrative in such a situation - because there would be nothing for me to mentally “point” to.

Phenomenal concepts, I argue (Chapter 3 of this dissertation), are indexicals that achieve reference not through an act of demonstration, but through the disposition of the subject to token that indexical concept when tokening the

phenomenal quality (in appropriate circumstances). As such, there is no need to postulate any representational mode of presentation for the indexical. Another way of putting this is that there is nothing available to the subject at a cognitively accessible or personal level through which the subject's concept is tied to its referent.

Demonstrative reference requires a mode of presentation. As many theorists have noted, there has to be some aspect of the thing referenced that enables reference to occur. For it can be rational to doubt that *this* is *that* even though 'this' and 'that' co-refer (Think, as examples, of demonstrations of the aft and stern of a ship, or of a sensation and a brain process (if some identity theory is true)).

However, it is never rational to believe that I am not me, or that today isn't today (barring worries about midnight striking mid-thought). This is because the reference of pure indexicals is *automatic*. It occurs at the sub-personal level and thus requires no mode of presentation to refer.

I have argued that phenomenal concepts are indexicals, and if this is so then there need be no mode of presentation at all for the phenomenal side of the psychophysical identity statement. When I imagine the world other than it is, I imagine that I am in a context in which the causal links between me and my phenomenal states are different.

I admit there is a strong inclination to think that phenomenal concepts are demonstratives. After all, we are obviously able to conceptualize phenomenal experience via demonstratives. I can, for example, think of *that* shade and refer to it later in thought.

One reason to doubt that phenomenal concepts are demonstratives is that we are often able to form phenomenal concepts of qualities we haven't experienced. [See chapter three].

Another reason is that our ability to demonstrate phenomenal qualities outstrips our ability to form phenomenal concepts. Consider two minimally distinct color tiles side by side. We notice that the shades are different and are thus able to demonstrate each shade (given distinctness in experience is sufficient for demonstration). However, we are utterly unable to distinguish them in thought. Given this, these demonstrative concepts cannot play the role typically assigned phenomenal concepts [see chapter one].

Note that this inability to possess distinct phenomenal concepts of the colors is exemplified in ways that go well beyond the inability to sort the colors in future encounters. Perhaps surprisingly, one can't even say, when looking at them, which is lighter, more saturated, or yellower, etc. It thus seems utterly unmotivated to postulate that the subject has acquired two distinct phenomenal concepts through these two demonstrations.

So how does it work then with non-demonstrative indexicals? Well, the reference is a function of causal facts that are unavailable at the personal level. I can, for instance, successfully refer to *here* and *now* without there being any modes of presentation. Same goes for phenomenal concepts.

So what happens when I doubt that *pain = C - fiber stimulation* is that I conceive that the primitive indexical that I token regularly when I feel pain is causally tied to something other than C-fiber stimulation. Thus there is a ready answer as to how *C - fiber stimulation = pain* can be a posteriori — an answer not available to those who take phenomenal concepts to be demonstratives (which plausibly involve non-discursive, but cognitively accessible, modes of presentation).

References

- [1] David M. Armstrong. Dispositions are causes. *Analysis*, 30:23–6, 1969.
- [2] Ned Block. Max Black’s objection to mind-body identity. In T. Alter and S. Walter, editors, *Phenomenal Concepts and Phenomenal Knowledge: New Essays on Consciousness and Physicalism*, pages 249–306. Oxford University Press, 2006.
- [3] David Chalmers. *The Conscious Mind*. Oxford University Press, 1996.
- [4] David Chalmers. The content and epistemology of phenomenal belief. In Q. Smith and A. Jokic, editors, *Consciousness: New Philosophical Perspectives*, pages 220–272. Oxford University Press, 2003.
- [5] David Chalmers. Phenomenal concepts and the explanatory gap. In T. Alter and S. Walter, editors, *Phenomenal Concepts and Phenomenal Knowledge: New Essays on Consciousness and Physicalism*, pages 207–219. Oxford University Press, 2006.
- [6] Noam Chomsky. Review of verbal behavior. *Language*, 35:26–58, 1959.
- [7] Donald Davidson. Thinking causes. In John Heil and Alfred Mele, editors, *Mental Causation*. Clarendon Press, 1993.
- [8] Daniel Dennett. *Consciousness Explained*. Little, Brown and Co., 1991.
- [9] Daniel Dennett. What Robomary knows. In T. Alter and S. Walter, editors, *Phenomenal Concepts and Phenomenal Knowledge: New Essays on Consciousness and Physicalism*, pages 15–31. Oxford University Press, 2006.
- [10] Justin Fisher. Why nothing mental is just in the head. *Noûs*, 41:318–34, 2007.
- [11] Jerry Fodor. *In Critical Condition*. MIT Press, 1998.
- [12] Christopher Hill. Imaginability, conceivability, possibility, and the mind-body problem. *Philosophical Studies*, 87:1:61–85, 1997.
- [13] Christopher Hill and Brian P. McLaughlin. There are fewer things than are dreamt of in Chalmers’s philosophy. *Philosophy and Phenomenological Research*, 59:445–454, 1999.
- [14] David Hume. *A Treatise of Human Nature*. Dover Publications, 2003.

- [15] Frank Jackson. Epiphenomenal qualia. *Philosophical Quarterly*, 32:127–136, 1982.
- [16] Frank Jackson. What Mary didn't know. *The Journal of Philosophy*, 83:285–291, 1986.
- [17] Frank Jackson. *From Metaphysics to Ethics: A Defense of Conceptual Analysis*. Oxford University Press, 1998.
- [18] David Kaplan. Demonstratives. In Joseph Almog, John Perry, and Howard Wettstein, editors, *Themes from Kaplan*. Oxford University Press, 1989.
- [19] Egon Peter Köster. Does olfactory memory depend on remembering odors? *Chemical Senses*, 30(s1):i236–i237, 2005.
- [20] Jaegwon Kim. Concepts of supervenience. *Philosophy and Phenomenological Research*, 45:153–76, 1984.
- [21] Saul Kripke. *Naming and Necessity*. Harvard University Press, 1980.
- [22] Janet Levin. What is a phenomenal concept? In T. Alter and S. Walter, editors, *Phenomenal Concepts and Phenomenal Knowledge: New Essays on Consciousness and Physicalism*, pages 87–110. Oxford University Press, 2006.
- [23] Joseph Levine. Materialism and qualia : The explanatory gap. *Pacific Philosophical Quarterly*, 64:354–361, 1983.
- [24] Joseph Levine. *Purple Haze: The Puzzle of Consciousness*. Oxford University Press, 2001.
- [25] Joseph Levine. Phenomenal concepts and the materialist constraint. In T. Alter and S. Walter, editors, *Phenomenal Concepts and Phenomenal Knowledge: New Essays on Consciousness and Physicalism*, pages 145–165. Oxford University Press, 2006.
- [26] David Lewis. What experience teaches. In William Lycan, editor, *Mind and Cognition*. Blackwell, 1990.
- [27] Brian Loar. Phenomenal states (revised version). In O. Flanagan N. Block and G. Gzeldere, editors, *The Nature of Consciousness*, pages 597–616. MIT Press, 1997.
- [28] Brian Loar. Phenomenal intentionality as the basis of mental content. In Martin Hahn and Bjørn Ramberg, editors, *Reflections and Replies: Essays on the Philosophy of Tyler Burge*. MIT / Bradford Books, 2003.
- [29] Colin McGinn. *The Problem of Consciousness*. Basil Blackwell, 1991.

- [30] Brian P. McLaughlin. Type epiphenomenalism, type dualism, and the causal priority of the physical. *Philosophical Perspectives*, 3:109–34, 1989.
- [31] Brian P. McLaughlin. Davidson’s response to epiphenomenalism. In John Heil and Alfred Mele, editors, *Mental Causation*. Clarendon Press, 1993.
- [32] Brian P. McLaughlin. Color, consciousness, and color consciousness. In Quintin Smith, editor, *New Essays on Consciousness*. Oxford University Press, 2003.
- [33] Thomas Nagel. What is it like to be a bat? *Philosophical Review*, 83:435–450, 1974.
- [34] Laurence Nemirow. Physicalism and the cognitive role of acquaintance. In William Lycan, editor, *Mind and Cognition*. Blackwell, 1990.
- [35] David Papineau. *Thinking about Consciousness*. Oxford University Press, 2002.
- [36] John Perry. *Knowledge, Possibility, and Consciousness*. MIT Press, 2001.
- [37] Elizabeth Prior, Robert Pargetter, and Frank Jackson. Three theses about dispositions. *American Philosophical Quarterly*, 19:251–57, 1982.
- [38] Willard van Orman Quine. *Ontological Relativity and Other Essays*. Columbia University Press, 1977.
- [39] John Searle. *The Rediscovery of the Mind*. MIT Press, 1992.
- [40] J.J.C. Smart. Sensations and brain processes. *The Philosophical Review*, 68:2:141–156, 1959.
- [41] Daniel Stoljar. Physicalism and phenomenal concepts. *Mind & Language*, 20:469–494, 2005.
- [42] Michael Tye. *Consciousness, Color, and Content*. MIT Press, 2000.
- [43] Michael Tye. A theory of phenomenal concepts. *Minds and Persons: Royal Institute of Philosophy Supplement*, 1:91–105, 2003.
- [44] Stephen White. Curse of the qualia. *Synthese*, 68:333–368, 1986.
- [45] Stephen White. Property dualism, phenomenal concepts, and the semantic premise. In T. Alter and S. Walter, editors, *Phenomenal Concepts and Phenomenal Knowledge: New Essays on Consciousness and Physicalism*, pages 210–248. Oxford University Press, 2006.

Curriculum Vita

Douglas Parvin

Education:

1994 A.B., Harvard College, Cambridge MA
2009 Ph.D., (Filing May 2009), Philosophy, Rutgers University, New
 Brunswick, NJ

Employment:

2008 – Augustana College, Instructor (Tenure Track), Rock Island, IL