


Consciousness and Intentionality

by

James R. John

B.A. Philosophy
New York University, 1998



ARCHIVES

SUBMITTED TO THE DEPARTMENT OF LINGUISTICS AND PHILOSOPHY IN
PARTIAL FULFILLMENT OF THE REQUIREMENTS FOR THE DEGREE OF

DOCTOR OF PHILOSOPHY IN PHILOSOPHY
AT THE
MASSACHUSETTS INSTITUTE OF TECHNOLOGY

SEPTEMBER 2004

© 2004 James R. John. All rights reserved.

The author hereby grants to MIT permission to reproduce
and to distribute publicly paper and electronic
copies of this thesis document in whole or in part.

Signature of Author: _____

A handwritten signature in black ink, appearing to be "James R. John", written over a horizontal line.

Department of Linguistics and Philosophy
July 23, 2004

Certified by: _____

A handwritten signature in black ink, appearing to be "Alex Byrne", written over a horizontal line.

Alex Byrne
Associate Professor of Philosophy
Thesis Supervisor

Accepted by: _____

Vann McGee
Professor of Philosophy
Chair, Committee on Graduate Studies



Consciousness and Intentionality

by

James R. John

Submitted to the Department of Linguistics and Philosophy
on July 16, 2004 in partial fulfillment of the
requirements for the Doctor of Philosophy in
Philosophy

ABSTRACT

This dissertation is about phenomenal consciousness, its relation to intentionality, and the relation of both to issues in the philosophy of perception.

My principal aim is (1) to defend an account of what it is for a perceptual experience to be phenomenally conscious and (2) to develop, within the terms set forth by this account, a particular theory of perceptual phenomenal consciousness.

Given the way these matters are usually understood, it probably is not obvious why I distinguish two philosophical tasks here. One might ask: “Isn’t defending an account of what it is for a perceptual experience to be phenomenally conscious the same thing as developing a particular theory of perceptual phenomenal consciousness?” I argue that it is not.

In addition to my principal aim, I have three subsidiary aims.

First, to shed some light on what it means for a perceptual experience to be an intentional mental event, one with representational content. Many philosophers regard the notion of perceptual intentionality as utterly unproblematic. Though I accept that experiences almost always have content, I subject this claim to more scrutiny than is usual.

Second, to go some way towards better understanding the relationship between perceptual phenomenal consciousness and perceptual intentionality. In particular, I examine recent attempts to explain the former in terms of the latter. My conclusion is that there can be no such explanation.

Finally, to show that, by improving our understanding of perceptual phenomenal consciousness, perceptual intentionality, and the relation between them, we can make headway on some very difficult problems in the philosophy of perception. I am especially interested in defending direct realism, the view that, in having perceptual experiences, subjects can be—and usually are—directly aware of material objects.

Thesis Supervisor: Alex Byrne

Title: Associate Professor of Philosophy

Acknowledgements

I'd like to begin by thanking my dissertation committee: Alex Byrne (chair), Alison Simmons, Bob Stalnaker, and Steve Yablo. Special thanks are due to Alex, with whom I've been working since my second year at MIT, and to Alison, for taking time out of her busy Harvard schedule to join my committee.

Though it has become something of a cliché to say this in MIT dissertation acknowledgements, it's true all the same: MIT is a wonderful place to study philosophy. I got a great deal of help and support from my fellow graduate students. Tyler Doggett and Andy Egan deserve first mention. They read multiple drafts of many of these chapters and were always willing to talk with me about my work. Selim Berker, Juan Comesaña (honorary MIT graduate student), Josh Flaherty, Liz Harman, Bernard Nickel, and Carolina Sartorio are also due special thanks. I'm grateful to Roxanne Fay, Alex Friedman, Patrick Hawley, Anthony Newman, Dilip Ninan, Chris Robichaud, Adina Roskies, Asta Sveinsdottir, and Eric Swanson. Each of them made helpful suggestions on how to improve earlier versions of my chapters. Finally, I want to thank an MIT professor not on my committee, Ned Hall, for his advice and encouragement throughout my final year at MIT.

I've been fortunate to be able to discuss my ideas with Ned Block, Mike Martin, and Susanna Siegel. I thank them for their helpful comments and criticisms. Thanks also go to Chrissy Graham, Mary Grenham, Stefanie Hanlon, Christopher Naylor, and Jen Purdy. We are very lucky at MIT to have such friendly and professional support staff.

I thank my parents, Conner and Janet John.

Finally, I owe my greatest debt of gratitude to Michela Ippolito. Without her love, patience, and understanding I never would have finished this.

This dissertation is dedicated to the memory of James P. Harford, who helped in so very many ways to make it possible.

Table of Contents

Introduction.....	9
Chapter 1: Phenomenal Consciousness.....	17
Chapter 2: Presentation, Location, and the Transparency Constraint.....	28
Chapter 3: The Confrontation Model.....	40
Chapter 4: Naïve Realism.....	51
Chapter 5: The Sense-Datum Theory.....	65
Chapter 6: Disjunctive Naïve Realism and the Trouble with Brute Confrontation.....	80
Chapter 7: The Possession Model.....	93
Chapter 8: Pure Representationism.....	111
Chapter 9: Impure Representationism.....	127
Chapter 10: Pure and Impure Phenomenism.....	149

Introduction

1. Aims of the Dissertation

This dissertation is about phenomenal consciousness, its relation to intentionality, and the relation of both to issues in the philosophy of perception.

My principal aim in writing this work has been (1) to defend an account of what it is for a perceptual experience to be phenomenally conscious and (2) to develop, within the terms set forth by this account, a particular theory of perceptual phenomenal consciousness.

Given the way these matters are usually understood, it probably is not obvious why I distinguish two philosophical tasks here. One might ask: “Isn’t defending an account of what it is for a perceptual experience to be phenomenally conscious the same thing as developing a particular theory of perceptual phenomenal consciousness?” I hope to convince you that it is not.

In addition to my principal aim, I have had three subsidiary aims.

First, to shed some light on what it means for a perceptual experience to be an intentional mental event, one with representational content. Many philosophers regard the notion of perceptual intentionality as utterly unproblematic. Though I accept that experiences almost always have content, I will subject this claim to more scrutiny than is usual.

Second, to go some way towards better understanding the relationship between perceptual phenomenal consciousness and perceptual intentionality. In particular, I will examine recent attempts to explain the former in terms of the latter. My conclusion will be that there can be no such explanation.

Finally, to show that, by improving our understanding of perceptual phenomenal consciousness, perceptual intentionality, and the relation between them, we can make headway

on some very difficult problems in the philosophy of perception. I am especially interested in defending direct realism, the view that, in having perceptual experiences, subjects can be—and usually are—directly aware of material objects.

2. The Aims Explained

Let me say a bit more by way of explanation about my dissertation’s principal and subsidiary aims. I begin with the principal aim, the one to do with phenomenal consciousness, then go on to discuss the subsidiary aims, the ones to do with, respectively, perceptual intentionality, perceptual phenomenal consciousness and perceptual intentionality, and the problems of perception.

Phenomenal Consciousness

Since the collapse of behaviorism in the 1950s and early 1960s, we have witnessed a surge of interest on the parts of both scientists and philosophers in phenomenal consciousness. Because of their abiding interest in the mind-body problem, philosophers have been most concerned with whether phenomenal consciousness can be reductively explained. They have asked what we can call the “reduction question”:

REDUCTION QUESTION—Can phenomenal consciousness be explained in terms derived solely from the physical sciences?

Physicalist reductionists answer in the affirmative, while dualist nonreductionists answer in the negative.¹

Though there is considerable diversity in the terms philosophers typically use to

¹ It will not matter for our purposes how, exactly, physicalist reductionism and dualist nonreductionism are best understood. Neither will it matter how the physical sciences are to be demarcated. But see Crane and Mellor (1990) for worries about whether the reduction question can be stated non-trivially.

formulate the reduction question—reductionists insist and nonreductionists deny that “sensations,” “conscious states,” “conscious events,” “qualia,” and “experiences” can be reductively explained²—most philosophers would agree, I think, on the following: if any mental events are genuinely phenomenally conscious, they are phenomenally conscious because they have phenomenal properties, and these phenomenal properties are what can or cannot be reductively explained.

This general, if usually tacit, agreement conceals something important. There has been and continues to be fundamental disagreement over what it is, ontologically speaking, for a mental event to have a phenomenal property. Philosophers have not reached a consensus regarding the proper way to answer what we can call the “structure question”:

STRUCTURE QUESTION—What is it for a mental event to have a phenomenal property?

As we will see, there are two historically important, mutually incompatible answers to the structure question. One of them, the answer given by the *Confrontation Model*, says that what it is for a mental event to have a phenomenal property is for the event’s subject to stand in a very special awareness relation—the “brute confrontation” relation—to instantiated properties. The other answer, the one given by the *Possession Model*, denies this, insisting that what it is for a mental event to have a phenomenal property is for the event to instantiate a non-confrontation-based property. For the Possession Model, these non-confrontation-based properties simply are the phenomenal properties.

Note that adopting either the Confrontation Model or the Possession Model fails to suffice for a full-blown theory of phenomenal consciousness. This is because each model furnishes a mere schema for a mental event’s being phenomenally conscious. We will discover

² See some of the older papers in Block, Flanagan, and Güzeldere (1997) for this plethora of terms.

that there are numerous applications of each model, some more plausible than others, which fill in the models' schemas to yield robust theories of phenomenal consciousness. For example, one application of the Confrontation Model might say that material objects are the bearers of the property instances with which the model says subjects are brutally confronted. Another might contend that sense-data, not material objects, bear these property instances.

Now it should be clear why I distinguished above between defending an account of what it is for a perceptual experience to be phenomenally conscious and developing a particular theory of perceptual phenomenal consciousness. The first philosophical task is completed by selecting and successfully arguing for either the Confrontation Model or the Possession Model. The second philosophical task is completed by selecting and successfully arguing for a particular application of the favored model. Answering the structure question, then, is a necessary, though often unacknowledged, prolegomenon to any theory of phenomenal consciousness.

Determining the right answer to the structure question—deciding whether to hold the Confrontation or the Possession Model—aids in answering the reduction question. Some contemporary philosophers, for example, are attracted to applications of the Possession Model that explain phenomenal consciousness in terms of intentionality, because they think doing so allows for a physicalist reductionist answer to the reduction question. (Of course, they would not express their theoretical preference in these terms.)

This dissertation does not attempt to answer the reduction question—though it does suggest that accepting the Confrontation Model's answer to the structure question forces a dualist nonreductionist answer to the reduction question. Rather, its concern is with the structure question. It defends an application of the Possession Model that we will call "pure phenomenism."

Perceptual Intentionality

Since Sellars (1956) and Anscombe (1965), it has become customary for philosophers to treat perceptual experiences as *intentional*, as having representational content.³ The fact that, in having an experience, the world can seem other than the way it is—think, for example, of the bent stick illusion—is usually taken as proof of the intentionality of perception. More recently, some philosophers have begun claiming that perceptual phenomenal consciousness and perceptual intentionality are closely related.

But too many contemporary philosophers proceed as if the answers to two questions are perfectly clear. The first question is philosophical: just what does it mean to say of perceptual experiences that they are intentional? After all, they seem radically different from occurrent thoughts, not to mention beliefs and desires—all paradigm examples of intentional mental events and states. The second question is historical: why was there a sudden shift (by the standards of the history of philosophy) from thinking of perceptual experiences as nonintentional (as, in the first instance, relations to non-propositional objects of perception) to thinking of them as intentional (as, in the first instance, relations to representational contents, usually conceived of as propositions)?

Ignoring the philosophical question is inadvisable, because there is an extremely intuitive line of thinking about perceptual experience that resists the thought that perception is intentional. Thus, we can allow that experiences are intentional only if we have a rebuttal of this line of thinking. This dissertation provides just that. Ignoring the historical question is inadvisable because, in the swift transition from the old non-representationalist regime to the new

³ This should be qualified. Many of the ancients, the medievals, and the early phenomenologists regarded perceptual experience as intentional, though it is questionable whether they conceived of intentionality in the way we do today. See chapter 1 of Robinson (1994) for discussion. In the modern continental tradition, we find

representationalist regime important insights—insights into the nature of phenomenal consciousness—were lost. We can regain them only by taking a second look at some of the recent history of the philosophy of perception. This dissertation does just that.

Perceptual Phenomenal Consciousness and Perceptual Intentionality

Some philosophers have recently begun arguing that there is an extremely tight connection between perceptual phenomenal consciousness and perceptual intentionality. One school (see, for example, Tye 1995 and 2000) seeks to explain the former in terms of the latter, while the other (see, for example, Siewert 1998) seeks to explain the latter in terms of the former.

The topic of the relation between the phenomenal and the intentional parts of our mental natures is enormous. Because of this, I will only discuss the project of explaining perceptual phenomenal consciousness in terms of perceptual intentionality. The four applications of the Possession Model we will consider below—what we will call “pure representationism,” “impure representationism,” and “impure phenomenism,” along with pure phenomenism—entail four different answers to the question whether an experience’s phenomenal consciousness can be explained in terms of its intentionality. The first three theories answer “yes.” This dissertation will argue by elimination for the fourth theory, pure phenomenism, which answers “no.”

The Problems of Perception

I think that by getting clearer on all of these issues we can make progress on some very difficult problems in the philosophy of perception. Here I will mention two.

One problem concerns the nature of perceptual experiences. What are they? What is

philosophers defending the intentionality of perception years before their analytic counterparts followed suit. See Hammond, Howarth, and Keat (1991) for discussion. My observation is intended to apply to analytic philosophy.

their structure, if any, like? Philosophers have failed to appreciate the extent to which views about what experiences are like—whether, for example, one adopts the “act-object” or the “adverbial” analysis of experience—are largely determined by views about phenomenal consciousness. I argue that the act-object analysis is motivated by antecedent commitment to the Confrontation Model, while the adverbial analysis is motivated by antecedent commitment to the Possession Model.

The other problem is one of the oldest in the history of the philosophy of perception: it is the problem of direct awareness.⁴ Here is one way of stating it. Consider the following two theses:

HALLUCINATION—Hallucinatory experiences—for example, when a subject has a perceptual experience as of an F object in her environment though there is no F object in her environment—are possible.

INDISTINGUISHABILITY—For every veridical experience a subject may have, there is a possible corresponding phenomenally indistinguishable hallucinatory experience.

These are attractive theses, and, though they are sometimes denied, most philosophers endorse them. But for centuries philosophers believed that the only acceptable theory of phenomenal consciousness that could accommodate and explain HALLUCINATION and INDISTINGUISHABILITY was one that entailed the denial of the following equally attractive thesis:

DIRECT AWARENESS—In having veridical experiences, subjects are directly aware of material objects.

Why did philosophers believe this? Because, to simplify enormously, they felt certain that HALLUCINATION and INDISTINGUISHABILITY, together with some plausible assumptions, force us to hold an application of the Confrontation Model—the notorious sense-datum theory—on which

⁴ In his recent eponymous book, Smith calls the problem of direct awareness the “problem of perception.” See Smith (2002).

it is impossible for subjects to be directly aware of material objects. If so, though, DIRECT AWARENESS is false. But if it is false, so is direct realism: the view that, in having perceptual experiences, subjects can be—and usually are—directly aware of material objects. This dissertation contends that all four applications of the Possession Model can accommodate and explain all three of these attractive theses.

3. A Look Ahead

To sum up, my dissertation defends the Possession Model and an application of it, pure phenomenism, that regards phenomenal consciousness as unexplainable in intentional terms. In doing so, it advances views on perceptual intentionality and some tough problems in the philosophy of perception.

Perhaps it goes without saying, but throughout this work I will concentrate on the case of phenomenally conscious perceptual—in particular, visual—experience. See chapter 1, though, for a defense of this limitation of scope.

Here is the plan. Chapter 1 introduces phenomenal consciousness and the philosophical problems it poses. Chapter 2 discusses the so-called “transparency” of experience. Chapter 3 presents the Confrontation Model. The Confrontation Model’s three applications—naïve realism, the sense-datum theory, and disjunctive naïve realism—are explained and criticized in the chapters 4, 5, and 6. Chapter 6 concludes with an argument against any application of the Confrontation Model. Chapter 7 presents the Possession Model. Pure and impure representationism are explained and criticized in chapters 8 and 9. Chapter 10 finishes things off with a critique of impure phenomenism and a defense of pure phenomenism.

Now it is time to take a closer look at our subject: phenomenal consciousness.

Chapter 1: Phenomenal Consciousness

1. Introduction

Chapter 1 begins by singling out from all of the many problems of consciousness the one with which we will be principally concerned. Next it explains the terminology we will use. It closes by restating the aims of the dissertation in the favored terminology.

2. The Problems of Consciousness

Perceptual experiences, bodily sensations, emotions and moods, acts of imagining, wondering, and even just thinking are all *conscious* mental events.⁵ Such mental events pose the philosophical problems of consciousness.

Why do I write ‘problems’ rather than ‘problem’? Because there are at least two problems of consciousness. Block distinguishes between *phenomenal consciousness* and *access consciousness* (see Block 1995).⁶ As Block’s discussion reveals, phenomenal consciousness and access consciousness pose very different philosophical problems.

A mental event is access conscious if it is poised for deployment in thought and the control of intentional action. The problem of access consciousness is the problem of explaining what it is for a mental event to be poised in this way.

The problem of phenomenal consciousness will be our concern. A mental event is phenomenally conscious if—to use the jargon popularized by Nagel (1974)—there is something

⁵ Some mental occurrences are best thought of as events, owing to their episodic nature, while others are best thought of as states, owing to their nonepisodic nature. I believe that phenomenally conscious mental occurrences should be treated as mental events, though I will not defend this belief. Thus, I will use the term ‘mental event’ throughout this dissertation. Nothing significant will turn on this choice of terminology.

⁶ Block’s distinction is related in various ways to the distinction in Chalmers (1996) between “phenomenal” and “psychological” concepts of mind.

it's like for a subject to undergo it. The problem of phenomenal consciousness is the problem of explaining how this can be: how can there be something it's like for a subject to see a Rothko or to feel a sharp, stinging pain?

But matters are more complicated than even Block's distinction leads us to believe, for the problem of phenomenal consciousness is itself comprised of two sub-problems. The first is the problem of *phenomenality*, and the second is the problem of *subjectivity*.

I said that the problem of phenomenal consciousness is the problem of explaining how there can be something it's like for a subject to undergo certain mental events. Notice that this breaks into two parts: the *what-it's-like* part and the *for-the-subject* part. The problem of explaining how there can be something it's like for a subject to undergo certain mental events is the problem of phenomenality. 'Phenomenality' refers to the qualitative or what-it's-like aspect of phenomenal consciousness.⁷ The problem of explaining how there can be something it's like for a subject to undergo certain mental events is the problem of subjectivity. 'Subjectivity' refers to the personal or for-the-subject aspect of phenomenal consciousness.

Solving the problem of phenomenality would mean accounting for the qualitative aspect of certain mental events. To see this, first look at your hand, then close your eyes and imagine it. There is an undeniable qualitative difference—a difference in what it's like—between the two mental events. What makes the visual experience qualitative in the distinctive way that it is? What makes the act of imagining qualitative in the distinctive way that it is? And what makes them both qualitative at all? This is the problem of phenomenality.

Solving the problem of subjectivity would mean accounting for the personal aspect of certain mental events. Consider again your visual experience of your hand and your act of

⁷ I picked up the term 'phenomenality' from discussions with Ned Block. I am unsure whether he uses it in exactly the way I use it here.

imagining it with your eyes closed. Though there is a dramatic qualitative difference between these states, they have at least one thing in common: they both have an important element of “for-me-ness” built in. When you look at your hand, there is something it is like *for you* to have the experience; when you close your eyes and imagine your hand, there is something it is like *for you* to do the imagining. The what-it’s-like part of each mental state is different, but the for-the-subject—the personal—part is the same. What is this for-me-ness that certain mental events have? This is the problem of subjectivity.⁸

It is impossible to imagine phenomenality and subjectivity coming apart. That is, it is impossible to imagine either (1) having a mental event that exhibits phenomenality (there is something it’s like to have it) without subjectivity (but there is no for-the-subject aspect to it) or (2) having a mental event that exhibits subjectivity (there is a for-the-subject aspect to it) without phenomenality (but there is nothing it’s like to have it). Try to imagine one without the other—you won’t be able to do it. Phenomenality and subjectivity, then, are a package deal: both are necessary for a mental event to count as phenomenally conscious.

If that is right, though, why is it necessary to distinguish between phenomenality and subjectivity? First, because there is a distinction to be made: there really are two conceptually separable aspects of phenomenal consciousness.⁹ Second, because doing so helps us to see more clearly exactly what is mysterious about phenomenal consciousness.

This dissertation is concerned exclusively with the problem of phenomenality. Unless otherwise indicated, then, when I mention phenomenal consciousness I mean to refer to

⁸ Levine (2001) also divides the problem of phenomenal consciousness into two sub-problems. In describing the problems the way I do, I borrow heavily from Levine. In particular, the ‘for-the-subject’ locution is his.

The problem of subjectivity is bound up with the question of what it is for a subject to have a perspective on the world in thought and experience. See Nagel (1986) on the link between the personal aspect of phenomenal consciousness and the perspectival character of mental life.

⁹ See p.215 of Byrne (2004) for doubts on this score.

phenomenality, the qualitative or what-it's-like aspect of phenomenally conscious mental events, rather than to subjectivity, the personal or for-the-subject aspect of such events. The theories we will consider below are all theories of phenomenality (though they might have things to say about subjectivity).

(Perhaps it goes without saying, but we are dealing with the *problems* of phenomenality and subjectivity because it is extremely difficult to understand what phenomenality and subjectivity are and how they can be parts of the world, especially as modern science describes it. Moreover, I do not intend my focus on phenomenality to suggest that I regard subjectivity as somehow easier to deal with. For what it is worth, the two problems strike me as equally difficult.)

Let me close this section with an important qualification. My attention throughout will be directed, with only a few exceptions, to the phenomenality of *visual* experience. Though I believe that everything I say about visual phenomenality applies to the phenomenality of experiences in the other sense modalities (and, indeed, to the phenomenality of nonperceptual phenomenally conscious mental events), I have elected not to discuss those other cases. I am aware that extending an account of visual phenomenality to, say, tactile phenomenality (not to mention nonperceptual phenomenality) is hazardous. But we can afford to set such worries aside. After all, even if what I have to say about visual phenomenality fails to apply happily to nonvisual phenomenality, there is still the substantive and hopefully interesting question whether my story about visual phenomenality is true.

Now it is time to introduce some important terminology. This terminology will help us to formulate more clearly the problem of phenomenality and the range of approaches to it that this dissertation will consider.

3. Phenomenal Properties and Representational Properties

We begin with ‘phenomenal property’ and proceed to ‘pure representational property’ and ‘impure representational property’.

Phenomenal Properties

Suppose you are having a veridical visual experience—call it “E”—of a dark green avocado against a neutral background. E is a mental event with phenomenality: there is something it’s like to have it.

What it’s like to have E—its phenomenality—is relevantly the same as what it’s like to have visual experiences of jade, pesto, and lawn clippings. E has a property in virtue of which this fact obtains, a *phenomenal property*—call it “G” (pronounced “big gee”). E has G in common with experiences of jade, pesto, and lawn clippings. What it’s like to have E—its phenomenality—is relevantly different from what it’s like to have visual experiences of stop signs, the clear sky, or ripe bananas. While E has G, experiences of stop signs have R, experiences of the clear sky have B, and experiences of ripe bananas have Y. R, B, and Y, like G, are phenomenal properties. A visual experience of an avocado and a ripe banana, viewed beside a stop sign with the clear sky in the background, would have all four phenomenal properties.

Thus, phenomenal properties are properties of phenomenally conscious mental events that type them according to what it’s like to have them—according, that is, to their phenomenality. They are properties of the mental states that have them. So G belongs to E, not to the avocado you see. The totality of an experience’s phenomenal properties at a time t constitutes the experience’s overall phenomenality—what it’s like, completely, to undergo it—at

t. To keep things simple, I will only consider E's G-ness and not any of its other phenomenal properties.

Note that I am using 'phenomenal property' in such a way that to say of a mental event that it has a phenomenal property is to say of it that it exhibits phenomenality. On this usage, explaining E's G-ness is explaining its phenomenality. E, of course, will have other properties, equally deserving of the title 'phenomenal property', in virtue of which it exhibits subjectivity. But, since I will have so little to say about subjectivity, I will speak of phenomenal properties only when I mean to speak of phenomenality and vice versa. From now on, then, having phenomenality *just is* having phenomenal properties. (Compare: having intentionality just is having representational properties.)

It should not be controversial that perceptual experiences have phenomenality. For all I have said, G might be almost *any* sort of property. For example, it might be the theoretically innocent property of being an experience as of something's looking green. (Theoretically innocent, because no one denies that E has *that* property.) The debates over phenomenality examined by this dissertation are debates over the nature of phenomenal properties, over what sort of property they are. To deny at the start that there are phenomenal properties, then, is just to deny the manifest fact of phenomenality. No one should do that.¹⁰

Some philosophers use, respectively, 'quale' and 'qualia' to mean what I mean by,

¹⁰ Two points.

First, I doubt that even Dennett, who famously implored us to "quine qualia" (Dennett 1988), would find anything objectionable in my talk of phenomenality. He is best read as urging us to reject a particular *theory* of phenomenal properties, not their *existence*.

Second, there are only two reasons I can see for rejecting my way of thinking of phenomenal properties. First, one might think I simply have not said enough to make the notion clear. There is little I can do to assuage this worry. Second, one might hold what Shoemaker calls the "Frege-Schlick" view (Shoemaker 1982) and tend towards a general skepticism of phenomenal notions. According to the Frege-Schlick view, intersubjective comparisons of phenomenal sameness and difference do not make sense. The Frege-Schlick view rests on verificationist assumptions; because of this, I will not discuss it. (However, for nonverificationist reasons to think that such comparisons do not make sense, see Stalnaker 1999).

respectively, ‘phenomenal property’ and ‘phenomenal properties.’ That is not how I will employ these terms. Phenomenal properties, as I understand them, constitute the phenomenality of phenomenally conscious mental events. All parties to the debate about phenomenal consciousness should agree that experiences have them. Qualia, on the other hand, are properties, which I will discuss in later chapters, posited by particular theories of phenomenality (to foreshadow: pure and impure phenomenism). As we will see, there are many philosophers who wish to reject these theories and, with them, qualia. So one can deny the existence of qualia (in my sense) without denying that there are phenomenal properties (in my sense).

Representational Properties

The next bit of terminology is the distinction between pure and impure representational properties.¹¹

Something has a pure representational property if and only if it has the property of representing a content. E, for example, plausibly represents the content that something is green. Thus, it has the pure representational property of representing that something is green.¹²

Something has an *impure representational property* if and only if it has the property of representing a content in a specific manner. Thus, something’s having an impure representational property can be factored into two parts: (1) the content it represents and (2) the specific manner in which it represents that content. Let me say a bit about this talk of the

¹¹ The labels ‘pure representational property’ and ‘impure representational property’ are Chalmers’s. See Chalmers (2004a).

¹² I assume that (most—see chapter 10) perceptual experiences represent contents and, consequently, have pure representational properties. For an explicit rejection of this assumption, see Alston (1998).

I also assume that the contents of experiences are propositions. For now, it will suffice to say just this about propositions: they are abstract objects that either are or determine possible worlds truth-conditions. For a Fregean treatment of perceptual content, see McDowell (1994). For a Russellian account, see Peacocke (1992). And for a possible worlds account, see Stalnaker (1998). I will have much more to say about the content of experience in later chapters.

specific manner in which something represents a content.

An assertion that something is green, the interpreted sentence ‘something is green,’ and a series of smoke signals indicating that something is green all have something in common: they all represent the content that something is green. Consequently, they all have the pure representational property of representing that something is green. But each of them represents that content in a different specific manner. The assertion represents it in an assertoric manner, the interpreted sentence represents it in a linguistic manner, and the series of smoke signals represents it in a, well, smoke signals manner.

Whenever something has the property of representing a content in a specific manner—whenever, that is, something has an impure representational property—it will have *nonrepresentational* properties that determine the specific manner in which it represents its content. Consider the assertion. It has nonrepresentational properties that make it an *assertion* that something is green rather than, say, a *supposition* that something is green. Talk of the specific manner in which something represents a content, then, is nothing more than talk of the nonrepresentational properties which make that something the type of representation—assertion, interpreted sentence, smoke signal, etc.—it is.

Applying all of this to E yields the following. E represents the content that something is green, so it has the pure representational property of representing that something is green. But E represents that content in a visual manner. To say this is to say that E has nonrepresentational properties which make it a visual representation rather than some other type of representation. Thus, E has the impure representational property of representing that something is green in a visual manner.¹³

¹³ As I employ the “pure” and “impure” terminology, the interpreted sentence ‘if something is green, then Abraham Lincoln is president’ has the impure representational property of representing that something green in a as-the-

Nothing can represent a content without representing it in a specific manner. This means that nothing can have a pure representational property without also having an impure representational property.

Anything with a pure representational property will have many corresponding impure representational properties. For example, E has the impure representational properties of representing that something is green in a perceptual manner, in a visual manner, in a visually attentive manner, etc. Each property in this list is more specific than the one preceding it. To represent a content in a visual manner is to represent it in a perceptual manner, but to represent it in a perceptual manner isn't necessarily to represent it in a visual manner, and the same goes for visual and visually attentive representation.

Philosophers often distinguish between the *representational content* of a mental event or state and the *attitude toward that content* which the event or state embodies. Take your perceptual belief, formed on the basis of having E, that something is green. It represents the content that something is green. As a belief, though, it is an instance of your taking the belief attitude towards that content.

This familiar distinction between content and attitude towards content is just a special case of the distinction between pure and impure representational properties. Your belief has the pure representational property of representing that something is green. This very abstract property is a matter of content. But your belief also has the impure representational property of representing that something is green in the belief manner. This much more specific property is a matter of attitude towards content.

Throughout this paper, when I say that a property is representational, I mean that it is

antecedent-of-a-conditional manner. I don't know whether Chalmers intends the terminology to apply this broadly. If not, though, no matter: nothing will turn on my broadening of it.

either a pure representational property or an impure representational property. When I say that a property is nonrepresentational, I mean that it is neither a pure representational property nor an impure representational property.

We can use the distinction between pure and impure representational properties to demonstrate the theoretical neutrality of my phenomenal property talk. For all I have said, phenomenal properties might be pure or impure representational properties. They might be nonrepresentational properties. At the moment, all of these options are open.¹⁴

4. Conclusion

I conclude this chapter with a restatement of the dissertation's aims, this time in our new terminology.

The principal aim is to defend pure phenomenism, an application of the Possession Model. Pure phenomenism, as we will see, identifies phenomenal properties with nonrepresentational properties of experience (the qualia mentioned above).

There are also three subsidiary aims.

The first is to better understand perceptual intentionality. What is it, as we can now say, for a perceptual experience to have a pure representational property? An impure representational property?

The second is to better understand the relationship between phenomenality and intentionality. I will argue that phenomenality cannot be explained in terms of either pure or impure representational properties.

¹⁴ Throughout this dissertation, I will assume that we seek a univocal account of phenomenal properties. That is, I assume that it is not the case that some phenomenal properties are, for example, representational while others are nonrepresentational. Though most philosophers make this assumption, I am unsure what justifies it. Block might be an exception: see the distinction between "mental paint" and "mental oil" on p.9 of his (forthcoming).

The third is to apply the dissertation's results to some of the problems of perception. We will see, among other things, that all four applications of the Possession Model—pure representationism, impure representationism, impure phenomenism, and pure phenomenism—count as forms of direct realism. Identifying phenomenal properties with either representational or nonrepresentational properties allows us to hold DIRECT AWARENESS without contradicting HALLUCINATION and INDISTINGUISHABILITY.

And now for some phenomenology.

Chapter 2: Presentation, Location, and the Transparency Constraint

1. Introduction

This chapter engages in a bit of phenomenology. After taking care of a preliminary terminological matter, I will do two things.

First, I will attempt to demonstrate that phenomenality is, like phenomenal consciousness more generally, complex. Just as phenomenal consciousness has two distinguishable aspects, phenomenality and subjectivity, phenomenality itself has two distinguishable aspects, what I will call “presentationality” and “locationality.”

Second, I will show that the notions of presentationality and locationality allow us to formulate an important phenomenological adequacy condition on theories of phenomenality, what I will call the “transparency constraint.”

2. Preliminary Terminological Matter

I have been speaking of your experience E’s representing the avocado as being green. But it will be better if we do not prejudge the question of which property E represents the avocado as having.

Let us speak, then, of E’s representing the avocado as being “g” (pronounced “little gee”). g is the property that seems instantiated by every part of the avocado’s surface, the property we are pre-theoretically inclined to say is green, the avocado’s color. Two things about g.

First, g is not G. G is a property of your experience, while g is a property of a material

object, the avocado you see.¹⁵

Second, though it might seem obvious that g is green, I will leave open the question of g's identity. I do this for two reasons. For one thing, not everyone agrees that g is green (see, for example, Shoemaker 1994 and Thau 2002). More importantly, the ensuing discussion will be easier to state if we remain noncommittal about g.

With this little matter out of the way, we are ready to proceed.

3. Presentation and Location

We have seen that the problem of phenomenality is the problem of explaining the qualitative or what-it's-like aspect of phenomenal consciousness. Take E. There is something it's like for you to have it. We draw attention to what it's like to have it when we characterize it as a G experience.

Now carefully introspect E and its G-ness. If you are worried about how to introspect E's G-ness, that's okay: just try to turn your attention to whatever it is that seems, well, G about it. (And if you think that has mainly, perhaps entirely, to do with the avocado's apparent g-ness, you are on to something—see the next section.) Introspecting in this way, I contend, reveals that E's phenomenality has two distinct phenomenological aspects, what we can call its "presentationality" and its "locationality."

Presentationality

Let us begin with presentationality. If you introspect E, you will find that there is something about what it's like to undergo it in virtue of which you are disposed to judge that in

¹⁵ Actually, there are theories on which (1) g is G and (2) E misrepresents the avocado as having a property that really belongs to your experience itself. Such views are implausible—see chapter 10. See also Shoemaker (1994).

having it E makes g present to your mind. This “something” is E’s *presentationality*. Its presentationality is one aspect of its phenomenality.

Take this talk of being disposed to judge that g is present to your mind with a few grains of salt. All I mean is that E is such that, as a matter of what it’s like to have it, you find it natural to characterize it using certain sorts of expressions. There is just something about what it’s like to have E in virtue of which you are spontaneously inclined to report things like the following: “When I have E, g is somehow *present to my mind, right there, before me*—all in a way it isn’t when, with eyes closed, I think thoughts about g things.” As a matter of what it’s like to have E, you find certain descriptions of it compelling. These descriptions evoke E’s presentationality.

To get a better grip on this, compare E with an occurrent thought you might form—call it “T”—about the avocado. T is a mental event directed at the avocado you see; it might even share representational content with E.¹⁶ Nevertheless, T lacks G, and, consequently, there is nothing about what it’s like to have T in virtue of which you are disposed to judge that it makes g present to your mind.

Two comments are in order.

First, I am *not* claiming that mental events with phenomenality stand in a special “presentation” relation to properties. Presentationality talk is not metaphysical talk. Rather, it aims to describe an aspect of phenomenality. It is phenomenological talk.

Of course, presentationality might be *explained* by appeal to a special relation between mental events with phenomenality and properties. For example, a theory might say that an experience has presentationality if and only if its subject is brutally confronted with a property, where the property the subject is brutally confronted with is the one her experience makes present

¹⁶ As an occurrent *thought* about the avocado, T is not—certainly need not be, anyway—accompanied by an act of imagining the avocado.

to her mind. (We will meet this relation of “brute confrontation” in later chapters.) But this invocation of the brute confrontation relation would furnish a candidate theoretical explanation of presentationality—not a definition of it.

Second, a mental event can have presentationality even though, in undergoing the event, its subject is not related to any property which could be said to make true the subject’s judgement that a property is present to her mind. Thus, that E has presentationality does not entail that E really does make a property “present to your mind.” I want to leave open whether, in having E, there is any property that we could describe as somehow present to your mind in having E. This point will not matter until chapter 10, but then it will be very important.

Locationality

Now let us turn to locationality. If you introspect E, you will find that there is something about what it’s like to undergo it in virtue of which you are disposed to judge that, for every property you take it to make present to your mind, E locates the property in your external, mind-independent environment.¹⁷ This “something” is E’s *locationality*. Its locationality is one aspect of its phenomenality.

Take this talk of being disposed to judge that E locates g in your external, mind-independent environment with a few grains of salt. All I mean is that E is such that, as a matter of what it’s like to have it, you find it natural to characterize it using certain sorts of expressions. There is just something about what it’s like to have E in virtue of which you are spontaneously inclined to report things like the following: “When I have E, g is present to my mind but also

¹⁷ I intend the sense of ‘locate’ on which a police officer might say that a witness located a suspect at a crime scene. According to the witness’s testimony, the suspect was at the crime scene. I do not mean the sense of ‘locate’ on which locating something is detecting it.

located by my experience in my environment. It's as if in having it E *says* that g is instantiated. It's as if there is a way the world is *according to* my experience." As a matter of what it's like to have E, you find certain descriptions of it compelling. These descriptions evoke E's locationality.

To get a better grip on this, compare E with an act of imagining the avocado—call it "I." I is a mental event directed at the avocado you see; it might even share representational content with E. Unlike your occurrent thought T, I has phenomenality. However, it lacks G, having instead a phenomenal property we can call "G*" (pronounced "big gee star"). That E has G while I has G* reflects an important difference in what it's like—an important difference in phenomenality—between your visual experience of the avocado and your act of imagining it. The two mental events are phenomenally similar (especially if you have a vivid imagination), but it would be very difficult to mistake one for the other.

Now, since I has G* instead of G, you are disposed to judge that in having it I presents a different property than the one you take E to present. Call this property "g*" (pronounced "little gee star"). So you are spontaneously inclined to say that E makes g present to your mind and that I makes, not g, but g* present to your mind. But E and I have something very important in common. There is something about both mental events in virtue of which you are disposed to say that they locate the properties they present in your external, mind-independent environment.

Three comments are in order.

First, I am *not* claiming that mental events with phenomenality stand in a special "location" relation to properties. Locationality talk is not metaphysical talk. Rather, it aims to describe an aspect of phenomenality. It is phenomenological talk.

Of course, locationality might be *explained* by appeal to a special relation between

mental events with phenomenality and properties. For example, a theory might say that an experience has locationality if and only if it has representational content according to which various properties are instantiated in its subject's environment, where the represented properties are the ones her experience locates. (We will come back to the connection between locationality and intentionality in later chapters.) But this invocation of the representation relation would furnish a candidate theoretical explanation of locationality—not a definition of it.

Second, mental events can have locationality even if neither their subjects nor anything in their subjects' environments has the properties the subjects take to be located. Even if E were a hallucination as of an avocado, you would still be disposed to judge that it locates g in your external, mind-independent environment.

Third, a mental event can have locationality even though, in undergoing the event, its subject is not related to any property which could be said to make true the subject's judgement that a property is located. Thus, that E has locationality does not entail that E really does "locate" a property in your external, mind-independent environment. I want to leave open whether, in having E, there is any property that we could describe as somehow located in having E. This point will not matter until chapter 10, but then it will be very important.

Why Care About Presentationality and Locationality?

To sum up, E is a G experience. Its phenomenality has two distinct phenomenological aspects, presentationality and locationality. You can tease them apart by carefully introspecting your experience.

I apologize for the complexity of my descriptions of presentationality and locationality—the prose is tortured, I know. In my defense, though, these are features of our conscious life that

are very difficult to put into words. Though I do not have—yet—explicit definitions of presentationality and locationality, I hope that my strategy of using the metaphors of presence and location help you to grok the phenomena.¹⁸

But why bother with it? For two reasons. First, because presentationality and locationality are real. Second, because highlighting them allows us a better understanding of the so-called “transparency” of phenomenally conscious experience—see the next section.

My talk of presentationality and locationality is intended as a theoretically neutral characterization of phenomenality. All parties to the consciousness debate, I contend, should agree that experiences have presentationality and locationality. Another advantage of distinguishing them is that doing so helps us recognize that a theory of phenomenality must explain two bits of phenomenology, not just one. As we will see, applications of the Confrontation Model tend to have a great deal to say about presentationality and little if anything to say about locationality. And one application of the Possession Model—pure representationism—has a great deal to say about locationality and little if anything to say about presentationality. An appreciation of presentationality and locationality helps in deciding between competing theories of phenomenality.

Some final points.

Just as it is impossible to imagine phenomenality and subjectivity coming apart, it is impossible to imagine presentationality and locationality coming apart. That is, it is impossible to imagine either (1) having a mental event which disposes you to judge that a property is present to your mind that doesn't also dispose you to judge that a property—the same property—is located or (2) having a mental event which disposes you to judge that a property is located that

¹⁸ I am not alone in finding talk of presence to mind helpful in characterizing the phenomenology of perceptual experience. Chapter 6 of Yolton (1996) traces the philosophical history of this metaphor. See chapter 4 of Arnauld

doesn't also dispose you to judge that a property—the same property—is presented. Try to imagine one without the other—you won't be able to do it. Presentationality and locationality, then, are a package deal: both are necessary for a mental event to count as having phenomenality.¹⁹

As a matter of what it's like to have it, you are disposed to judge that E presents and locates g. I take it that what subjects are most basically aware of in having their perceptual experiences are the properties they judge to be presented and located by their experiences.

I also think that the properties subjects are most basically aware of in having their experiences are the properties subjects' experiences represent their environments as instantiating. So when you judge that E presents and locates g, the property your judgement concerns is one your experience represents your environment as instantiating.

E represents the avocado as being g. g is also the property you judge E to present and locate. Your occurrent thought T about the avocado represents the avocado as being g too. But, since T lacks G, you do not judge it to present and locate g. It does have phenomenality, though. In particular, it has "inner saying" phenomenality: in thinking T, it is for you as if a quiet voice utters the sentence you would use to express T. T, then, has "inner saying" phenomenal properties; thus, you are spontaneously inclined to say that it presents and locates "inner saying"-associated properties. Occurrent thoughts, then, are not such as to dispose their subjects to judge that they present and locate the same properties they represent. This explains why many feel the temptation to deny that occurrent thoughts have phenomenality. It is easy to feel that the

(1990) for a fascinating discussion of its role in early modern philosophy.

¹⁹ Some might grant the impossibility of imagining locationality without presentationality, while insisting that it is possible to imagine presentationality without locationality. They might even cite bodily sensations such as pains or orgasms as examples of mental events with presentationality but no locationality.

I deny that these mental events are genuine cases of presentationality without locationality. Pains dispose you to judge that they locate damage properties in your body. Orgasms dispose you to judge that they locate, well, orgasm properties in your body.

phenomenal properties of occurrent thoughts belong instead to “inner saying” experiences that merely accompany the thoughts.

4. The Transparency Constraint

We have seen that phenomenality—the qualitative or what-it’s-like aspect of phenomenal consciousness—itsself has two distinguishable aspects: presentationality and locationality. In this section, I will show that we can put the notions of presentationality and locationality to work in formulating a perspicacious statement of the alleged *transparency* of phenomenally conscious experience.

Harman draws our attention to the transparency of experience in the following passage:

When Eloise sees a tree before her, the colors she experiences are all experienced as features of the tree and its surroundings. None of them are [sic] experienced as intrinsic features of her experience. Nor does she experience any features of anything as intrinsic features of her experience. And that is true of you too. . . . When you see a tree, you do not experience any features as intrinsic features of your experience. Look at a tree and try to turn your attention to intrinsic features of your visual experience. I predict you will find that the only features there to turn your attention to will be features of the presented tree, including relational features of the tree “from here” (1990, p.667).

Tye continues in the same vein:

Focus your attention on the scene before your eyes and on how things *look* to you. You see various objects; and you see these objects by seeing their facing surfaces. . . . Intuitively, the surfaces you directly see are publicly observable physical surfaces. . . . In seeing these surfaces, you are immediately and directly aware of a whole host of qualities. You may not be able to name or describe these qualities but they look to you to qualify the surfaces (2002, p.138, emphasis Tye’s).

The following comes later:

None of the qualities of which you are directly aware in seeing the various surfaces *look* [sic] to you to be qualities of your experience. You do not experience any of these qualities *as* qualities of your experience (ibid., emphasis Tye’s).

And finally:

The qualities you experience are the ones the surfaces apparently have. Your experience is thus transparent to you. When you try to focus upon it, you ‘see’ right through it, as it were, to the things apparently outside and their apparent qualities (p.139).

Thus, Harman and Tye say that when we introspect our perceptual experiences we find that the only properties we “experience” or are “immediately and directly aware of” are properties that look to belong to things in our external, mind-independent environments—not to our experiences. Notice that there are really three claims here.²⁰

The first claim is that, in having our experiences, we “experience” or are “immediately and directly aware of” properties. When they characterize our experiences in this way, using terms like ‘experience’ and ‘direct awareness,’ they are attempting to describe what it is about phenomenally conscious experience in virtue of which subjects are disposed to judge that in having their experiences properties are made present to their minds. Thus, Harman and Tye are drawing our attention to the presentational aspect of phenomenality (though they would not put their point this way).

The second claim is that the properties we “experience” or are “immediately and directly aware of” in having our experiences look to belong to objects in our external, mind-independent environments—not to our experiences. When they speak of these properties looking to belong to things in the external world, they are attempting to describe what it is about phenomenally conscious experience in virtue of which subjects are disposed to judge that the properties they take to be present to their minds in having their experiences are located by their experiences in the external, mind-independent world. Thus, Harman and Tye are drawing our attention to the locational aspect of phenomenality (though they would not put their point this way).

The third and final claim is that when we introspect our perceptual experiences, we find

²⁰ Moore first brought transparency to the attention of philosophers in the analytic tradition. He famously wrote: “And, in general, that which makes the sensation of blue a mental fact seems to escape us: it seems, if I may use a metaphor, to be transparent—we look through it and see nothing but the blue...” (1993, p.37). He also wrote: “When we try to introspect the sensation of blue, all we can see is the blue: the other element is as if it were diaphanous” (ibid, p.41).

that the only properties we are aware of are the properties we “experience” or are “immediately and directly aware of,” the properties which look to belong to things in our immediate mind-independent environments—not to our experiences. When you introspect E you discover that the only properties you are aware of in having your experience are the properties you judge to be presented and located. The crucial point is that since you are spontaneously inclined to report that these properties are properties of objects in your environment, none of them is such as to strike you as belonging *to your experience itself*.

These three claims come together to yield a statement of what it is for a perceptual experience to be transparent. A perceptual experience is transparent if and only if (1) it has presentationality and locationality and (2) introspection of it reveals that, in having the experience, the only properties the subject is aware of are the properties she is disposed to judge are presented and located.

Note that, necessarily, experiences have phenomenal properties if and only if they are transparent. Try imagining an experience phenomenally the same as E that is not transparent—you won’t be able to. And try imagining an experience that is transparency-wise the same as E but which lacks phenomenal properties—again, you won’t be able to. This shows that phenomenality and transparency are of a piece: mental events with phenomenality are transparent and vice versa.

Using the notion of transparency we can formulate the following phenomenological adequacy condition on theories of phenomenality:

TRANSPARENCY CONSTRAINT—If all else is equal, reject theories of phenomenality that conflict with transparency.

There are four ways for a theory of phenomenality to conflict with transparency and, thus, to violate the transparency constraint.

Three are insignificant. First, the theory could entail that experiences lack presentationality. Second, the theory could entail that experiences lack locationality. Third, the theory could entail that introspection finds awareness of properties other than the ones subjects judge to be presented and located. I doubt any theory of phenomenality has any of these consequences. To entail any of these things would be to deny the very feel of phenomenally conscious experience. No theory I know of has these consequences.

The fourth way *is* significant. A theory of phenomenality could entail that the properties subjects judge to be presented and located by their experiences are not instantiated where their experiences locate them. Take the sense-datum theory, an application of the Confrontation Model. It says that the properties subjects are most basically aware of in having their experiences are instantiated by sense-data, not material objects in subjects' external, mind-independent environments (see chapter 5 for more on sense-data). But, in virtue of their experiences' locationality, subjects are disposed to judge that these properties belong out there in the external world—not to sense-data. If the sense-datum theory is right, though, these judgements are mistaken. If possible, we should eschew theories of phenomenality that entail this sort of error.²¹

5. Conclusion

I have tried to show that phenomenality is complex, with distinguishable presentational and locational aspects, and I have offered an account of the transparency of experience.

²¹ Thus, we should deny that G is among the properties you judge to be presented and located. If it were, you would wrongly report it to belong to something in your environment. Indeed, I deny that subjects are aware of their experiences' phenomenal properties in having those experiences.

But how can I believe that while also holding the claim—assumed by our earlier comparisons of E with occurrent thoughts and acts of imagining—that phenomenal properties can be objects of introspective awareness? See Chapter 10. There I will show that there is no tension between thinking (1) that phenomenal properties are introspectible and (2) that subjects are not aware of phenomenal properties in having their experiences.

Chapter 3: The Confrontation Model

1. Introduction

In the introduction, I claimed that there are two historically important, mutually incompatible answers to the structure question, the question of what it is for a mental event to have a phenomenal property: the Confrontation Model and the Possession Model.

This chapter introduces the Confrontation Model.

2. The Confrontation Model and the Possession Model

Since at least the beginning of the early modern period, philosophers have argued over whether direct realism is true. Direct realism, recall, is the claim that, in having perceptual experiences, subjects can be—and usually are—directly aware of material objects. This debate has taken many turns. There have been times, such as the early modern period, when indirect realists were in the majority. There have been other times—we are living in such a time—when direct realists were in the majority.

This conflict can seem strange to us. Indeed, there are contemporary philosophers who, assuming that direct realism is surely correct, have wondered what all of the fuss could have been about. These philosophers tend to dismiss whatever philosophical pressure might have been felt in the past towards indirect realism as proceeding from a few egregious logical mistakes.

This interpretation of the history is wrong, and I will explain why in the next few chapters. But even if I am right, a question remains: what were the debates over direct realism debates about?

I think I can answer this question. I contend that the fight over direct realism is, in large part, a fight over how to think about phenomenality, a fight, in other words, over the best way to explain phenomenal properties. My view is that, running beneath the surface of the direct realism debate, is a conflict over which of two big-picture models of phenomenality to accept. These models are the Confrontation Model and the Possession Model.

These models have shaped the way philosophers have thought about perception. But for all of their influence, they have gone largely unnoticed. In fact, only now are philosophers of mind and perception beginning to see their outlines.

The two models are mutually incompatible. If one is true, the other must be false, and vice versa. As we will see, the Possession Model is, for the most part, simply the denial of the Confrontation Model. Thus, there is the possibility that, between the two of them, they do not exhaust the logical space of views. However, since each model is only a bare bones schema, in need of filling in to yield a full-fledged theory of phenomenality, I suspect that every theory of phenomenality that has ever been proposed counts an application of one of the two models.

My picture of things looks like this. Philosophers have found themselves fighting over direct realism and related issues such as the nature of perceptual experience. Arguments for and against various positions on these issues are advanced. To us, with centuries of hindsight, some of these positions seem bizarre and unmotivated. But this appearance is dispelled once we recognize the debate as conditioned by a subterranean, largely unarticulated conflict between two conceptions of phenomenality. Then we can see that philosophers found certain dialectical moves natural to make—and totally ignored others—because their philosophical intuitions were guided by tacit commitment to either the Confrontation Model or the Possession Model.

If successful, this dissertation will accomplish three things.

First, it will achieve the aims set out in the introduction.

Second, it will furnish something like a rational reconstruction of an important debate in the recent history of the philosophy perception, the debate over direct realism, and show how it has been all along mostly about phenomenality.

Finally, it will serve as a kind of philosophical therapy for its author. My MIT second-year paper was a defense of the sense-datum theory, a form of indirect realism. Everyone, my committee and my fellow students alike, thought I was crazy. I was astonished by their reaction: it seemed *obvious* to me that the sense-datum theory—perhaps not any particular philosopher's version of it, but some form of it—just had to be correct. I now realize that an unacknowledged commitment to the Confrontation Model was behind my certainty. In arguing against the Confrontation Model and its applications, I am showing the fly the way out of the fly bottle.

Which reminds me of something very important. I will devote a considerable amount of space to discussing views, the Confrontation Model and its applications, that might strike you as highly implausible. Give or take a little, I don't really disagree. But there are two very important reasons to take the views seriously. For one thing, there are quite a few living philosophers who endorse the Confrontation Model and some form of one or other of its applications. For another (and this is the most important reason), the Confrontation Model, through its naïve realist application, conditions almost all of our really deep intuitions about phenomenality and perceptual experience. We are, in a way that I will explain below, *committed* to it—even when we explicitly reject it. Bringing it to light and exploring its manifold weaknesses (and even, believe it or not, some of its strengths) will put us in a better position to assess our intuitions about phenomenality and perceptual experience.

(Note that the Confrontation Model and the Possession Model pertain exclusively to

phenomenality, not subjectivity. This was not made perfectly clear in the introduction. The two are not unrelated to theories of subjectivity, as one's choice of either model will have ramifications for what sort of view of subjectivity one will take. But they are independent.)

This chapter discusses the Confrontation Model. The next three chapters take on the model's three main applications: naïve realism, the sense-datum theory, and disjunctive naïve realism. In subsequent chapters, we turn to the Possession Model and its applications.

3. What is the Confrontation Model?

The Confrontation Model is easy to state, harder to explain. Let ' Φ ' stand for phenomenal properties like G and let ' ϕ ' stand for properties like g. This gives us:

CONFRONTATION MODEL—What it is for a mental event to have a phenomenal property Φ is for the event's subject to be brutally confronted with a local instance of ϕ .

According to the Confrontation Model, if you take a subject, a local instantiation of a property ϕ , and place them in the brute confrontation relation, the subject will undergo a mental event with a phenomenal property Φ .

To understand this, we need to understand two things. First, the schematic nature of the claim. Second, the contribution that talk of subjects, brute confrontation, and local instantiations of properties like g—what we can call, generically, “ ϕ -properties”—make to the claim's meaning. Before getting to these tasks, though, a brief word on the motivation for the Confrontation Model.

Why Hold the Confrontation Model?

Reflect again on E, your visual experience of the avocado. E exhibits phenomenality; it

has the phenomenal property G. E's phenomenality has two distinguishable aspects, presentationality and locationality. E is transparent: introspection of it reveals that, in having E, the only properties you are aware of are the properties you are disposed to judge are presented and located.

This shows us something very important. In having E, the only property you seem to yourself to be aware of is g, the property you are spontaneously inclined to report as presented and located. In particular, you do not seem to yourself to be aware of any property you would describe as a property of E. Neither do you seem to yourself to be aware of E itself or any mental mechanism that could serve to make you aware of g. You simply open your eyes and find g, with nothing else turning up upon introspection. This is why 'transparency' is such a good name for this feature of experience: in having E, your mind strikes you as being totally open to a range of properties.

The Confrontation Model is motivated by the thought that *all there is* to E's having G is your coming into some sort of brute contact with—confronting—an instantiation of g. Why are you disposed to judge that in having it E makes a property present to your mind? Because, according to the model, you stand in a relation to this property—the brute confrontation relation—which is such that when you enter into it with a property the property strikes you as being present to your mind, right there, before you. Why are disposed to judge that E locates the property it presents in your external, mind-independent environment? Because, according to the model, you can be brutally confronted, not with the property simpliciter, but only with an instantiation of it in your local environment. And why is this presented and located property the only property you are aware of in having E? Because, according to the model, there is nothing else—nothing hidden from introspective view, so to speak—involved in E's having G.

Now we can try to understand better exactly what the Confrontation Model says.

The Confrontation Schema

The Confrontation Model, as I noted in the introduction, furnishes a mere schema for a mental event's having phenomenality. It answers the structure question, but that is all: it does not give us, by itself, a theory of phenomenality. It is only a schema, because it leaves unspecified (1) what instantiates the ϕ -properties and (2) which properties the ϕ -properties are. Filling in those details, as we will see, does yield a theory of phenomenal properties, of what phenomenality is. There are three principal applications of the Confrontation Model: naïve realism, the sense-datum theory, and disjunctive naïve realism. Each application has a different account of (1) and (2). We will meet them in the next three chapters.

(What of the Confrontation Model's explicit reference to the *subjects* of mental events, a reference we will not find in the Possession Model? Unlike the nature of the ϕ -properties and the things that instantiate them, I do not count talk of the subject as something in need of filling in. The reason for this is that the nature of subjects is not something on which any application of the model makes the nature of phenomenality turn. This is highlighted by the fact that all applications of the model are free to give the same account of what it is to be the subject of a mental event.)

Subjects, Brute Confrontation, and Local Property Instances

Now that we understand the schematic nature of the Confrontation Model's claim, we can turn to the question of understanding what it means.

As for the reference to subjects, I will have nothing to say about this. I have no theory of

what it is to be the subject of a mental event. But this will not matter for our purposes, as understanding the model will not depend on understanding any such theory.

What about brute confrontation? The brute confrontation relation is a two-place relation between a subject and an instantiated property.

Brute confrontation is *brute* because it is primitive and unanalyzable. In particular, brute confrontation is not analyzable in representational terms. That is, for any instantiated property F, a subject's being brutally confronted with F is not to be analyzed as the subject's being in a mental state with the content that ...F....

That the brute confrontation relation is, in this way, nonrepresentational does not rule out the possibility that mental events with phenomenality have representational content. The Confrontation Model leaves space for its applications to claim that, though being brutally confronted with an instance of a ϕ -property is not itself representational, the mental event that involves its subject's standing in this relation to a property instance can have content. I will have more to say about this in subsequent chapters.

Note, also, that the bruteness of brute confrontation does not render applications of the Confrontation Model incompatible with the causal theory of perception. Indeed, one application of the model, the sense-datum theory, is almost always conjoined with the causal theory. I will have more to say about this as well in subsequent chapters.

Brute confrontation is *confrontation* because it is metaphysically direct: it is impossible to be brutally confronted with a property instance by being brutally confronted with a distinct property instance. (See below for a brief discussion of the connection between brute confrontation and the relation that is often called "direct acquaintance.")

The brute confrontation relation is supposed to be cognitive, cognitive in the sense that,

in standing in it to a property instance, the mind somehow grasps the property in a way that puts the subject in a position to have propositional knowledge.

The Confrontation Model makes a subject's standing in the brute confrontation relation to a local instance of a ϕ -property necessary and sufficient for phenomenality. What is meant by "a local instance of a ϕ -property"?

Clarification of two parts of this question must await later chapters: what the ϕ -properties are and what sort of object instantiates them. But why does the Confrontation Model require that the ϕ -properties be locally instantiated, and what does it mean to be brutally confronted with a property instance?

As I hinted above, locationality is the answer. In having E, you are disposed to judge that E locates g in your external, mind-independent environment. You are spontaneously inclined to report that g is instantiated in your environment. Thus, the very phenomenology of E suggests that g is instantiated nearby.

Two points. First, I do not know how to understand the locality requirement. It is probably relative. Even if I were to gain the ability to see objects as far away as the next galaxy, my experiences of such objects would still possess locationality and the properties they would strike me as presenting and locating, if any, would seem located, well, *near me*. Second, when I say that brute confrontation is a relation between subjects and property instances, perhaps what I ought to say is that it is a relation between subjects and *tropes*, which are said to be concrete universals. If you find this sort of talk helpful, feel free to substitute 'trope' everywhere I write 'property instance' or one of its cognates.

As evidence that friends of the Confrontation Model really did accept its local instantiation condition on confronted ϕ -properties, consider the following famous passage from

Price:

When I see a tomato there is much that I can doubt. I can doubt whether there is a tomato that I am seeing, and not a cleverly painted piece of wax. I can doubt whether there is any material thing there at all. Perhaps what I took for a tomato was really a reflection; perhaps I am even the victim of some hallucination. One thing however I cannot doubt: that there exists a red patch of a round and somewhat bulgy shape, standing out from a background of other colour-patches, and having a certain visual depth, and that this whole field of colour is directly present to my consciousness. What the red patch is, whether it is physical or psychical or neither, are questions that we may doubt about. But that something is red and round then and there I cannot doubt (1932, p.3).

Here is what I think is going on in the passage. Price is, for argument's sake, simply stipulating that he knows that he is having an R experience as of a ripe tomato. Now Price, I submit, holds the Confrontation Model, so he knows that, since he is having an R experience, he is brutally confronted with r (he, of course, just takes r to be red). But by the local instantiation condition, he knows that if he is brutally confronted with r, something—maybe not a tomato, but *something*—must locally instantiate r. This r something is what is “directly present to my consciousness.” Thus, since he stipulated that he knew he was having an R experience and since he knows that the Confrontation Model is true, he cannot doubt that “something is red and round” and that it is “directly present to my consciousness.” This is typical Confrontation Model-type reasoning.

Price could just as easily have said the following. He could have split the Confrontation Model into two claims. The first—call it “CONFRONT”—holding that what it is for a mental event to have Φ is for the event's subject to be brutally confronted with an *instantiated or uninstantiated* ϕ -property. This claim requires that the subject be brutally confronted with the property itself, not an instance of it. The second—call it “INSTANTIATE”—holding that, if a ϕ -property is brutally confronted, it must be locally instantiated: there are no cases of brute confrontation with locally uninstantiated properties—universals, if you like. Then Price could have reasoned as follows: “I hold CONFRONT but not INSTANTIATE. Thus, since I know

that I'm having an R experience, I know, by CONFRONT, that I'm brutally confronted by r. But, since I reject INSTANTIATE, I deny that it follows that anything *is* r. I'm brutally confronted by a locally uninstantiated property."

But Price did not reason this way. Neither did most of the other historical proponents of the Confrontation Model.²² The reason why, I claim, is that, impressed by the locationality of experience, they found the idea that subjects could brutally confront locally uninstantiated properties bizarre at best, unintelligible at worst. Think about it: what would it be to be visually aware of red—not an instance of red, not a red thing, just red. It will help bring out the negative reaction I am fishing for to put it this way: what would it be to be visually aware of a universal?²³

(I spoke in the previous two paragraphs as if Price could articulate things in my way and as if he "knew" the truth of the Confrontation Model. Price neither could articulate things in my way—simply because he didn't think of things in this way—nor did he know that the Confrontation Model is true. What I am doing is rationally reconstructing the thinking that, I claim, underlies Price's view.)

The passage from Price hints at something it is important to bring out about the Confrontation Model. The Confrontation Model motivates a certain way of thinking of the structure of perceptual experiences. The idea is this. Perceptual experiences are mental events constituted by direct awareness relations between subjects and objects. Which objects? The objects that instantiate the ϕ -properties subjects brutally confront in having their experiences. This is, in short, the so-called act-object analysis of experience: experiences are mental acts in

²² According to Baldwin (1990, p.234), Stout and Dawes Hicks, two of Moore's early critics, probably reasoned in the alternative way.

²³ I hope you find this as odd-sounding as I do. You might not, though, insisting that visual awareness of red is just having a red-representing experience. Fair enough. But see the chapters below on the Possession Model.

which subjects become directly aware of “objects of perception.” What is the direct awareness relation between subjects and objects of perception? Different philosophers have given different treatments of it. Usually, though, it is considered to be the “direct acquaintance” relation discussed in Russell (1912).

So commitment to the Confrontation Model underlies commitment to the act-object analysis of experience. It follows from what I have said that my brute confrontation relation is not the direct acquaintance relation that Russell and others spoke of.²⁴ But this additional complication is well worth making, for talk of brute confrontation makes clear why philosophers believed in direct acquaintance. Something, according to the Confrontation Model, must instantiate the confronted ϕ -properties. Whatever those objects are, be they material objects or sense-data, we surely stand in a very intimate relation to them: this relation is direct acquaintance. My brute confrontation relation, then, is the truly fundamental relation. Commitment to it, even if unarticulated, precedes commitment to direct acquaintance.

My brief remarks here on the Confrontation Model are intended only as a sort of distillation of a trend of thought that I detect in the recent history of the philosophy of perception. I will expand on this point in the next three chapters.

3. Conclusion

Now that I have introduced the Confrontation Model, it is time to have a look at its three applications: naïve realism, the sense-datum theory, and disjunctive naïve realism.

²⁴ Note that direct acquaintance was usually regarded as a relation between subjects and objects, not subjects and properties or property instances. However, Russell sometimes speculated that subjects could be directly acquainted with properties.

Chapter 4: Naïve Realism

The soul, he said, is composed
Of the external world.

—Wallace Stevens, from “Anecdote of Men by the
Thousand”

1. Introduction

The plan for the next few chapters is as follows.

In this chapter, I will present naïve realism and make three criticisms of it. Two of these are new. The other, that naïve realism cannot accommodate the fact of nonveridical perceptual experience, is very old, at least as old as theorizing about perception itself. This criticism leads directly to the sense-datum theory.

In the next chapter, I will present the sense-datum theory. It arises out of an attempt to maintain the Confrontation Model in the face of the problems posed by nonveridical experience. Though I will present difficulties for it, I will also offer a qualified defense of the sense-datum theory. I will argue that contemporary attacks on the theory—in particular, those which allege that the theory rests on a fallacy—are unfounded.

After that, in chapter 6, I will do two things. First, I will present and criticize disjunctive naïve realism. Its roots are like the sense-datum theory’s in that it is the product of an attempt to salvage the Confrontation Model in the face of its other applications’ failings. Second, I will mount an argument that applies, not only to disjunctive naïve realism, but to any possible application of the Confrontation Model. I contend that the very notion of brute confrontation, the relation which lies at the heart of the model’s way of explaining phenomenality, is incoherent.

2. Naïve Realism Introduced

As an application of the Confrontation Model, naïve realism fills in the blanks we found in the statement of the model. It provides an account of what sort of objects instantiate the confronted ϕ -properties and of what the ϕ -properties are.

Naïve realism adds to the Confrontation Model the following two theses:

MATERIAL OBJECTS—Material objects, and nothing else, instantiate ϕ -properties.

INTRINSICNESS—The ϕ -properties are intrinsic properties of material objects.

Take E. By MATERIAL OBJECTS, the object that instantiates g, the ϕ -property you brutally confront, is the avocado you see.²⁵ By INTRINSICNESS, g is an intrinsic property of the avocado's surface.²⁶

Naïve realism, then, says that what it is for E to have G is for you to be brutally confronted with g, where g is instantiated by the avocado you see and is an intrinsic property of its surface. Thus, it entails that G is the property of being an experience whose subject is brutally confronted with g instantiated by a material object. (Note: that you see the particular avocado you in fact see is not necessary for E to have G—any g-instantiating material object would do. So if you see a g avocado and I see a g cucumber, both of our experiences have G.)

According to naïve realism, perceptual experiences are mental events constituted by direct awareness relations between subjects and objects. Which objects? The material objects that instantiate the ϕ -properties subjects brutally confront in having their experiences. Thus,

²⁵ Naïve realism is free to be quite liberal about which things count as material objects. It can allow that clouds, rainbows, even shadows and the sky all count as material objects.

²⁶ Usually, this is conjoined with the claim that g is green. Naïve realism tends towards color—and, in general, secondary quality—primitivism.

naïve realism accepts the act-object analysis of experience.²⁷

3. Naïve Realism Explained

Why hold naïve realism? Because of transparency. Why are you disposed to judge that E makes g present to your mind? Because you stand in the intimate brute confrontation relation to an instantiation of it. Why are you disposed to judge that E locates g in your external, mind-independent environment? Because it is instantiated by the avocado that is before your very eyes. Why, upon introspecting E, do you find that you are aware of nothing but the avocado's g? Because this awareness of g—brute confrontation with g—that you enjoy in having E is all there is to E's having G. For the naïve realist, what you see, so to speak, in experience is exactly what you get. Nothing is hidden from our introspective view. It tolerates no mental mechanisms serving to facilitate perfect, glorious, phenomenality-generating contact with the material world. Phenomenality, according to naïve realism, is almost magical.²⁸

But only *almost* magical. Naïve realism is not obviously incompatible with the view that perception is a causal process. What's almost magical, according to naïve realism, is the proximal phenomenal-experiential effect of the causal chain leading from the distal avocado and the light reflecting off of its surface to you; however, there is nothing magical about the causal chain itself. The upshot of that process is that you undergo a mental event in which you enter into the brute confrontation relation with a local property instance.

It might help to put naïve realism as follows. It says that whatever is phenomenal about

²⁷ Most philosophers equate the act-object analysis with the sense-datum theory. But that is a mistake. Many different theories, naïve realism among them, subscribe to the analysis.

²⁸ Chalmers has recently claimed that a view very much like what I am calling naïve realism is "edenic": it is the view of phenomenality and perceptual experience we intuitively held in our pre-theoretical perceptual Garden of Eden, where we lived before eating the fruit of the tree of Knowledge of Nonveridical Experience. See Chalmers (2004b).

E—whatever is the ultimate source of your experience’s phenomenality—is quite literally on the surface of the avocado. According to it, then, the instantiation of *g* is what is “lights on!” about E. All you have to do, to have a *G* experience is to open your eyes, thereby creating the conditions for brute confrontation to occur, and look in the direction of a *g* thing. (Of course, you must also be awake.) Thus, naïve realism sees phenomenal properties of experience as only derivatively phenomenal. Instantiations of ϕ -properties are what are basically phenomenal in the world. Phenomenal properties *count as phenomenal* only because they are defined in terms of relations between subjects and instantiations of ϕ -properties. This is true of every application of the Confrontation Model, only non-naïve realist applications will have different accounts of what instantiates the ϕ -properties (and of what the ϕ -properties are). Indeed, for naïve realism, the perceptible properties of material objects are phenomenal precisely in the sense that sense-datum theorists regarded sense-datum properties as being.

Though naïve realism is not incompatible with thinking of perception as a causal process, it probably is incompatible with thinking that E has representational content. But even if it is compatible with this, it is definitely incompatible, in virtue of the nonrepresentational nature of brute confrontation, with thinking that content plays a role in determining phenomenality.

According to naïve realism, E does not represent *g* as instantiated in your immediate environment—it *presents* *g* as instantiated in your immediate environment.

I contend that naïve realism is our pre-theoretical stance to the world. By this, I mean that, because of the very phenomenology of phenomenally conscious perceptual experience (in particular, because of its transparency), we are disposed to comport ourselves to the world *as if* naïve realism were true. It is the attitude we take to the world in our various goings on. I call this a “stance” and not a folk theory because I do not think that naïve realism has the status of a

folk theory. A folk theory, I take it, is a set of more or less prescientific opinions held by ordinary subjects. Folk theories are such that their believers make explicit appeal to them in explaining various features of the world. Naïve realism is not like this. No one ever appeals to naïve realism in explaining why things are the way they are.²⁹ Moreover, unlike, say, folk psychology, we give it up relatively easily.

But even when we give it up—after learning about nonveridical experience or vision science—two things are worth noting. First, we feel a definite sense of astonishment—sometimes even anxiety—when we learn that naïve realism is false. Suppose I take a school child and induce a hallucinatory experience in him as of a red spot on the wall before him. How, he might ask in his bewilderment, could there not be something red before him? How could whatever is “reddish” (Levine 2001) or “red-feeling” (Byrne and Hilbert 1997) about his experience not be solely a matter of there being something red before him? (We will come back to such questions in the next chapter.) Second, even when the astonishment fades, we can sometimes find ourselves regarding the views of experience we accept once we have discarded naïve realism as somehow guilty of *bad faith*. These theories seem, in an important sense, to *lie* about experience; they turn their backs, so to speak, on the way phenomenally conscious experience presents itself to us as a radical phenomenal openness to the perceived world.³⁰

²⁹ Perhaps I should qualify this. I suspect that if anthropologists were to find a remote group of people who had no knowledge of perception science, they would discover that members of this group, if given a brief tutorial on how to use terms like ‘experience’ and ‘phenomenal property,’ give a naïve realist answer to this question: “Why does your experience have G?” I also suspect that those members of the group who had come to grips with the fact of nonveridical perception would hold something like the sense-datum theory.

Of course, I am making the condescending assumption that there isn’t a thriving native school of philosophers of perception in these peoples’ midst. So perhaps I should qualify my assertion by appealing to children. At all events, what is definitely true is that naïve realist explanations are not offered with anything like the frequency with which folk belief-desire explanations are offered.

³⁰ Two things.

First, see Valberg (1992) for an illuminating discussion of bad faith and our relation to the phenomenology of our own experiences. I don’t think such issues have been discussed enough by analytic philosophers, though this is beginning to change. See McDowell (1994), Johnston (2004), and Chalmers (2004b).

You might be thinking something like this: “Waddya mean by all this? Naïve realism is crazy! So how could we regard it this way?” I agree (sort of): given all we know, naïve realism is extremely implausible. But that does not change the fact that we do think about experience and phenomenality in the ways I have just described. To see this, note that, though no one endorses naïve realism exactly the way I have presented it here, all of us tend to think of it as a sort of regulative ideal. It is the theory we wish were true. That this is so explains why we want so badly to believe direct realism, the view that subjects can be—and usually are—directly aware of material objects. It is because phenomenally conscious experience presents itself to us in naïve realism’s terms that denials of indirect realism seem so, well, offensive. Though Hume thought that the “slightest philosophy” (1977, p.104) sufficed to overthrow naïve realism, he struggled mightily with the sometimes anxiety-provoking conclusion of his indirect realist reasoning.³¹

Does all of this mean that denying naïve realism commits us to an error theory about experience and phenomenality? I’m not sure. We would be committed to an error theory if our folk theory of experience and phenomenality were a naïve realist theory (just as we are committed to a moral error theory if moral realism—our folk moral theory—is false), but naïve realism is not our folk theory of experience and phenomenality. In fact, I do not think we have a folk theory of experience and phenomenality. So, from this perspective, it looks as if talk of an

Second, though I have intentionally limited the scope of my discussion to visual experience, I have said that I believe it to apply to the full range of phenomenally conscious mental events. The example of moods, though, gives me pause. Is there such a thing as naïve realism about elation or depression?

³¹ Has anyone ever held naïve realism?

Well, no one has held it in the terms I’ve stated it in. But perhaps some philosophers have endorsed it. Aristotle’s theory of perception in *De Anima* can be read as a form of naïve realism (see Aristotle 1984). According to Smith (2002), the existentialist phenomenologists, particularly Heidegger, seem to have held it. Johnston (2004) advances a modified form of it. See chapter 6 for comments on the views of self-professed contemporary naïve realists like M.G.F. Martin. Recent advocates of the “theory of appearing” like Harold Langsam and William Alston also seem to hold naïve realism (see Langsam 1995 and Alston 1998 and 1999). Indeed, Langsam claims that he is a theorist of appearing in order to defend naïve realism.

error theory is out of place. However, the way experience presents itself to us is such that putting it into words yields naïve realism. If naïve realism is false, then, to this extent, there is something false or misleading about the way experience presents itself. I am unresolved.

Before we turn to the problems facing naïve realism, I should note one more thing about it. *Phenomenal internalism* is the claim that phenomenal properties supervene on subjects' intrinsic properties: for any two possible subjects S_1 and S_2 , if S_1 and S_2 are intrinsic duplicates, then their experiences have the same phenomenal properties. *Phenomenal externalism* denies this: it is possible for two intrinsic duplicates to have experiences with different phenomenal properties. Naïve realism entails phenomenal externalism. If your intrinsic duplicate is not brutally confronted with g instantiated by a material object, then his experience lacks G even though yours has it. I will have much more to say about phenomenal internalism and externalism in subsequent chapters.

4. Naïve Realism Criticized

Now it is time to criticize naïve realism.

(Remember: at the end of the chapter 6, I will present a criticism that applies to all three applications of the Confrontation Model.)

Locationality?

Though it seems fairly clear that naïve realism has the theoretical resources to explain presentationality, it is less clear that it has the resources to explain locationality. Why does the mere fact of g 's being instantiated by the avocado you see account for locationality, the aspect of what it's like to have E in virtue of which you are disposed to judge that g is located in your

environment?

This is also a good place to mention that it is very difficult to see what story naïve realists can tell to explain which “local” property instances are local enough to enter into the brute confrontation relation to subjects. The story cannot be that property instances belonging to *visible* objects are local enough. The reason why not is that which objects count as visible is a matter of which objects cause *phenomenally conscious* experiences enabling subjects to distinguish them from their surroundings. Thus, to know whether an object is visible we need to know the very thing we don’t know: when a property instance is local enough to bring about a phenomenality-having experience.³²

Only One Brute Confrontation Relation?

Naïve realism’s use of the brute confrontation relation faces a dilemma. (Note: this problem also affects disjunctive naïve realism.)

Applications of the Confrontation Model have two parts. They explain phenomenality in terms of (1) brute confrontation with a property instance and (2) the particular property instanced. That is, you do not have a G experience unless you have brute confrontation with an instantiation of g. If you have brute confrontation with a different property—say r—you don’t get a G experience. If you have g but no brute confrontation with it, you don’t get a G experience. Question: how does the phenomenality-determining work get divided between (1) and (2)?

As applied to naïve realism, there are at least two options.

On the first option, there is only one brute confrontation relation. Standing in this

³² Naïve realism also owes us an account of why subjects can be brutally confronted with properties like g but not with natural or social kind properties or properties like the property of being your favorite avocado.

relation to a property instantiation is sufficient for phenomenality. But, once the phenomenal lights are on, so to speak, phenomenal differences are determined by differences in the confronted ϕ -properties. Compare E with a visual experience E* of an avocado that is exactly the same shape as the avocado you see in having E only painted r. E and E* both involve your standing in the same brute confrontation relation to a property instance. But E relates you to g while E* relates you r. The phenomenal difference between E and E* is entirely due to the difference in confronted property. This is simply naïve realism as I presented it at the beginning of this chapter. It is the standard way of applying the Confrontation Model.³³

The trouble with this option is that it fails to explain phenomenal differences between perceptual experiences in different sense-modalities. The trouble is generated by the possibility that experiences in different sense-modalities might nevertheless be experiences of the same property. Suppose you were to handle the avocado, having a tactile experience E** of its shape. It is plausible that both E and E** are experiences of the avocado's shape. But E and E**, we can assume, have no phenomenal properties in common. However, if there is only one brute confrontation relation and if both experiences are experiences of one and the same property of the avocado, then E and E** ought to have at least one phenomenal property in common. They don't. Call this the "modality problem."

One way to respond to this would be to deny that experiences in any two different sense modalities can be experiences of one and the same property. This might be the case, though I doubt it. Besides, if there is just one—*only* one—property of material objects that can be perceived in more than one sense modality, then the first division of phenomenal labor strategy fails.

³³ Moore's version of the sense-datum theory endorsed this strategy. He thought that all there was to his "sensation of blue" was the blue itself. See Moore (1993).

The best response takes us to the other option for dividing the phenomenal labor between brute confrontation and confronted properties. On the second option, there is more than one brute confrontation relation, one corresponding to each sense-modality. Standing in the *visual* brute confrontation relation to a property instance suffices for visual phenomenality, and within the visual modality phenomenal differences are determined by differences in confronted ϕ -properties. Now we can allow that E and E** are both experiences of the avocado's shape (so they involve brute confrontation with the same ϕ -property) without having to say that E and E** share at least one phenomenal property. Why not? Because this option allows us to account for the phenomenal difference between the visual experience of the avocado's shape and the tactile experience of the avocado's shape by appeal to different brute confrontation relations: *visual* brute confrontation in the former case, *tactile* in the latter. No modality problem.

But this option faces a different problem. In virtue of what could there be different confrontation relations? The only plausible answer is that visual experiences have special properties in virtue of which they involve visual brute confrontation, while tactile experiences have distinct special properties in virtue of which they involve tactile brute confrontation. Call the modality-specific properties that determine, for each modality, the type of brute confrontation it involves "specifying properties." This specifying property strategy is fair enough. But if it is right, there is more to E's having G than brute confrontation with a confronted ϕ -property. A necessary ingredient in determining whether E has G rather than R or some other phenomenal property will be which specifying properties it has. That is, to know whether a given mental event has G, we need to know which confronted ϕ -property is experienced, whether the event's subject is brutally confronted with the property, and, to determine whether the brute confrontation relation is of the right type, which specifying properties the event has. This strategy essentially

abandons the transparency-based motivation for the Confrontation Model. Call this the “specifying property problem.”

So the naïve realist appeal to brute confrontation faces a dilemma. Either there is only one brute confrontation relation, in which case we face the modality problem, or there are many modality-specific brute confrontation relations, in which case we face the specifying property problem. What this boils down to is this: if we take the first strategy, we are unable to give a plausible account of phenomenal differences across sense modalities, while taking the second strategy means giving up on the primary motivation for naïve realism—indeed, the Confrontation Model itself.³⁴

The Problem of Direct Awareness

Recall the three theses we discussed at the end of the introduction:

HALLUCINATION—Hallucinatory experiences—for example, when a subject has a perceptual experience as of an F object in her environment though there is no F object in her environment—are possible.

INDISTINGUISHABILITY—For every veridical experience a subject may have, there is a possible corresponding phenomenally indistinguishable hallucinatory experience.

DIRECT AWARENESS—In having veridical experiences, subjects are directly aware of material objects.

Naïve realism offers a beautiful explanation of the truth of DIRECT AWARENESS, for it predicts that subjects are directly aware of the objects, material objects, that instantiate the confronted ϕ -properties. The trouble with naïve realism—the one philosophers have grappled with for a long, long time—is that, given a very plausible assumption, it is incompatible with HALLUCINATION

³⁴ We will see in a later chapter that pure representationism, an application of the Possession Model recently defended by Michael Thau, faces a version of the same dilemma.

and INDISTINGUISHABILITY.

To see this, let us consider what we can call, following the discussion in Valberg (1992), a “Valberg case.” There you are, enjoying E, your G experience of the avocado you see, when, in an instant, Descartes’s Evil Demon destroys the avocado while nevertheless bringing it about that everything continues to appear visually the same to you. You are totally unaware of any change. Call the experience you undergo beginning with the switch “H.” E is a veridical experience. H is a hallucinatory experience. The transition from E to H is phenomenally seamless, so E and H are phenomenally indistinguishable.³⁵ Finally, let us suppose that the avocado was the only material object in your vicinity.

E and H are precisely the sort of veridical experience-hallucinatory experience pair that HALLUCINATION and INDISTINGUISHABILITY allow for. What this Valberg case does is to make the case for the following assumption: for most such phenomenally indistinguishable pairs, each member of the pair has the same phenomenal properties as the other. If this assumption is right, E and H are both G experiences. We should accept that the assumption is right because the best explanation of the phenomenal indistinguishability between E and H is that they are phenomenally the same.³⁶

But now the trouble for naïve realism is obvious. It says that G is the property of being a mental event whose subject is brutally confronted with a local instance of g. But by hypothesis

³⁵ Nothing will turn on whether E and H are numerically distinct experiences or merely time-slices (or whatever) of a single ongoing experience, one which goes from being veridical to being hallucinatory with the Evil Demon’s action.

³⁶ Two points.

First, it does not follow from the fact that E and H are phenomenally indistinguishable that they are phenomenally identical. The non-transitivity of phenomenal indistinguishability shows this. Suppose there are three color chips, each one a lighter shade of blue than the next one. You cannot distinguish your experience of chip 1 from your experience of chip 2. Neither can you distinguish your experience of chip 2 from your experience of chip 3. But you can distinguish your experience of chip 1 from your experience of chip 3. If phenomenal indistinguishability entailed phenomenal identity, we would get the result that your experiences of chips 1 and 2

there are no local instances of *g* when you undergo *H*: the avocado, now destroyed, was the only *g* material object in your vicinity. Indeed, if we were to suppose that there were no other material objects near you when you had *E*, naïve realism predicts, not only that *H* will lack *G*, but that it will lack phenomenal properties altogether: *E will have no phenomenality*. And given the necessary link between phenomenality and subjectivity, naïve realism predicts that, as soon as the avocado is destroyed, there will no longer be anything it is like for you to have your experience. (In fact, given naïve realism, it is questionable whether *H* even counts as a perceptual experience.) Things ought to go visually blank. Not dark, but blank—totally blind, the way things are for you visually now *behind your back*.

Thus, naïve realism, together with an extremely plausible assumption, is incompatible with HALLUCINATION and INDISTINGUISHABILITY. This, the problem of direct awareness, is a good reason to reject it.^{37,38}

have the same phenomenal properties, that your experiences of chips 2 and 3 have the same phenomenal properties, but that your experiences of chips 1 and 3 have different phenomenal properties.

Second, some disjunctivists deny the plausible assumption. I'll have more to say about them shortly.

³⁷ We could have made the same anti-naïve realism point using the phenomenon of illusion rather than hallucination. (Surely you could have an experience phenomenally indistinguishable from *E*—and, by our plausible assumption, a *G* experience—even though the material object you see (say, a pomegranate) is not *g*.) I prefer using hallucination cases. They are much easier to state and, as the example of the totally visually blank experience dramatizes, more gripping. Philosophers have been right to obsess over hallucination. Their possibility is a disturbing and important challenge to our thinking about perception and phenomenality.

Some philosophers might wish to deny the possibility of *E*-to-*H*-type transitions like our Valberg case. Such a denial is implausible. Indeed, surely we know enough about the brain to know that, through sophisticated enough extensions of current techniques, we could bring about an experience phenomenally indistinguishable from *E*.

³⁸ Theorists of appearing say something roughly like the following. What it is for *E* to have *G* is for you to stand in the primitive and unanalyzable three-place “appearing” relation to an object and a property. In having *E*, the avocado appears *g* to you. Appearing is *not* to be analyzed in representational terms. They take my talk of brute confrontation with property instances and direct awareness of objects and collapse them into a single three-place relation. The only difference is that they deny that the confronted properties must be instantiated. (Why? To allow that if in a case of illusion the avocado appears red to you, you can stand in the appearing relation to the avocado and the locally uninstantiated property *r*.)

Do they have a good account of hallucination? No. Here's why.

First, Alston would argue (see his 1999) that in *H* a “mental image” appears *g* to you. This is either the sense-datum theory (see the next chapter) or Russell's intermittently held view that we perceive our own brains in experience. Neither is very attractive.

Second, Langsam, whose view differs slightly from the way I just characterized the theory of appearing, would argue (see his 1995) that in *H* a “region of space” appears *g* to you. I am not sure what this means. Either the

5. Conclusion

One way to respond to this problem would be to reject the Confrontation Model and the way of thinking that leads to naïve realism. But that isn't quite what philosophers did. Historically, most philosophers responded, not by abandoning the Confrontation Model, but by embracing a non-naïve realist application of it: the sense-datum theory.

To that theory we now turn.

region of space in question instantiates *g* or it doesn't. If it does, why are others unable to see its instantiation? You can, after all, be the only one in a room hallucinating an avocado. If it doesn't, then what is it for a region of space to appear to have a locally uninstantiated property? The only plausible way of understanding this, I submit, is by seeing it as your experience's having the (false) representational content that something is *g*. But this pushes Langsam in the direction of some application of the Possession Model.

I conclude that theory of appearing versions of naïve realism cannot account for phenomenally indistinguishable hallucination.

Chapter 5: The Sense-Datum Theory

1. Introduction

This chapter discusses the sense-datum theory.

The sense-datum theory has enjoyed many defenders. Many of the early moderns, Locke and Hume especially, are best read as holding a position very much like it.³⁹ It was the dominant theory of perception among early analytic philosophers.⁴⁰

But more recently the sense-datum theory has come under strong attack and, consequently, has had fewer proponents. Nevertheless, it does have its contemporary supporters. Robinson (1994) is a vigorous defense of the sense-datum theory. The first chapter of Peacocke (1983) endorses it. Foster (2000) presents arguments for a variation on it called the “sense qualia theory.” Even Frank Jackson has held it (see Jackson 1977), though he has since repudiated it.

2. The Sense-Datum Theory Introduced

As an application of the Confrontation Model, the sense-datum theory fills in the blanks we found in the statement of the model. It provides an account of what sort of objects instantiate the confronted ϕ -properties and of what the ϕ -properties are.

The sense-datum theory replaces naïve realism’s MATERIAL OBJECT and INTRINSICNESS with the following two additions to the Confrontation Model:

SENSE-DATA—Sense-data, and nothing else, instantiate ϕ -properties.

³⁹ See Locke (1964) and Hume (1977). But see Yolton (1984) for doubts about whether my interpretation of the early moderns is accurate.

⁴⁰ See, among many others, Moore (1993), Russell (1912), Broad (1923), and Price (1932). See chapter 2 of Robinson (1994) and chapter 8 of Baldwin (1990) for commentary. Swartz (1965) is a collection of the classic early- and mid-twentieth century analytic papers for and against the sense-datum theory.

INTRINSICNESS SD—The ϕ -properties are intrinsic properties of sense-data.

Take E. By SENSE-DATA, the object that instantiates g, the ϕ -property you brutally confront, is a sense-datum, an immaterial mental object.⁴¹ By INTRINSICNESS SD, g is an intrinsic property of the sense-datum's surface.⁴²

The sense-datum theory, then, says that what it is for E to have G is for you to be brutally confronted with g, where g is instantiated by a sense-datum and is an intrinsic property of its surface. Thus, it entails that G is the property of being an experience whose subject is brutally confronted with g instantiated by a sense-datum.

According to the sense-datum theory, perceptual experiences are mental events constituted by direct awareness relations between subjects and objects. Which objects? The sense-data that instantiate the ϕ -properties subjects brutally confront in having their experiences. So the sense-datum theory accepts—indeed, is the most historically important instance of—the act-object analysis of experience.

3. The Sense-Datum Theory Explained

Let me say a bit more by way of explanation about the sense-datum theory.

Though some sense-datum theorists have had doubts about this, sense-data, as I noted above, were usually thought of as immaterial mental objects.⁴³ Also, they were usually taken to be private, in that only you have access to the sense-data your experiences relate you to.

Just as with naïve realism, the sense-datum theory claims that the objects instantiating the

⁴¹ Some sense-datum theorists thought of sense-data as two-dimensional, but sometimes—as in Peacocke (1983)—they thought of them as three-dimensional.

⁴² Sense-datum theorists believe that sense-data instantiate a range of both “primary” ϕ -properties and “secondary” ϕ -properties. This poses a host of difficult problems, but I will not discuss them.

⁴³ Moore, in particular, was given to waffling over whether sense-data are mental. See chapter 8 of Baldwin (1990) for a thorough discussion of Moore's thought on this question.

confronted ϕ -properties are the objects of which subjects are directly aware in having their experiences. This has an important consequence: it means that the sense-datum theory is incompatible with direct realism. I will have much more to say about this below in section 5.

The sense-datum theory is compatible with the causal theory of perception. Indeed, Grice's classic defense of that theory is couched in terms of the sense-datum theory (see Grice 1961).

The sense-datum theory is also compatible with the view that experiences have representational content. Perhaps more to the point, it is compatible with the claim that experiences have content about the external world. Contemporary philosophers have a curious tendency to deny this.⁴⁴ Perhaps that is because the sense-datum theory makes crucial appeal to the nonrepresentational brute confrontation relation. But a sense-datum theorist is free to say the following: "Perceptual experiences have representational content. Whatever theory of intentionality you like explains this. Brute confrontation with instantiations of ϕ -properties (and the consequent direct awareness of sense-data) is just there to add phenomenality. Take that away, and you're left with a bare representational mental event, something a bit like an occurrent thought. With it, though, you've got a phenomenally conscious perceptual experience." This line of thought goes well with a view on which instantiations of ϕ -properties by sense-data causally co-vary with instantiations of distal visible properties by material objects in subjects' environments, so that experiences, with their ϕ -propertied sense-data, serve as proximal "signs" of distal environmental features.⁴⁵

The sense-datum theory is very plausibly taken to entail phenomenal internalism: if you

⁴⁴ It is easy to read parts of Byrne (2001) this way.

⁴⁵ At all events, sense-datum theorists ought to deny that experiences have content about sense-data. If that content were connected to phenomenology in any way, it would get the phenomenology of experience wrong: experiences are of our physical environments, not sense-data.

are directly aware of a *g* sense-datum, so is your intrinsic duplicate.

4. The Sense-Datum Theory (Sort of) Defended

Why hold the sense-datum theory? Because of the “argument from hallucination.”⁴⁶ In this section, I will present a version of the argument, showing how the sense-datum theory arises out of naïve realism’s failure to handle hallucination. Then I will exonerate the argument from hallucination—well, the first stage of it anyway—of the charge that it commits an egregious “sense-datum fallacy.”

The Argument

The argument from hallucination comes in two stages, the first argues for direct awareness of sense-data in hallucination and the second argues for direct awareness of sense-data even in veridical experience.

The first stage goes as follows. Recall our Valberg case. Your veridical experience *E* and your hallucination *H* are phenomenally indistinguishable. As we saw, the best explanation of this is that *E* and *H* have the same phenomenal properties. But, given the Confrontation Model, this means that in having *H* you brutally confront an instance of *g*. There is no material object in your immediate vicinity. But *something* must be there to bear the confronted ϕ -properties. That something must be an immaterial object—call this thing an “idea,” “sensus,”

⁴⁶ Well, *most* sense-datum theorists have relied on some form of the argument from hallucination. But when Jackson held the sense-datum theory, he unequivocally rejected it. Writing of the argument from hallucination (and some related arguments) he avowed that “I think these arguments prove nothing” (1977, p.107).

Margaret Wilson has argued in more than one place that the motivation for early modern versions of the sense-datum theory has less to do with phenomenological concerns and more to do with worries about the relationship between “ideas” and the material world, especially as conceived by corpuscularian science. See the essays, especially the first, in her (1999).

or “sense-datum.” This sense-datum, then, is what you are directly aware of in having H.⁴⁷

The second stage—what Smith calls the “generalizing step” (Smith 2002, p.26)—extends this result to the case of veridical experience. Suppose we accept the conclusion of the first stage of the argument: in having H, you are directly aware of a g sense-datum. But what is going on when you undergo E? We seem to have three options:

Option 1: In having E, you are directly aware of a g sense-datum, and that is the only object you are directly aware of.

Option 2: In having E, you are directly aware of both a g sense-datum and a g material object, the avocado.

Option 3: In having E, you are directly aware of a g material object, the avocado, and that is the only object you are directly aware of.

Let us rule out Option 2. This is justified for two reasons. First, no one ever seems to have held it. Second, for reasons that I will not discuss here, it is implausible. (See Johnston 2004 for some of the reasons.) But why prefer Option 1 to Option 3?

Sense-datum theorists should not defend Option 1 by claiming that since E and H are phenomenally indistinguishable, they must be identical direct awareness-wise. For as Austin famously protested:

But if we are prepared to admit that there may be, even that there are, some cases in which ‘delusive and veridical perceptions’ really are indistinguishable, does this admission require us to drag in, or even let in sense-data? No. For even if we were to make the prior admission...that in the “abnormal” cases we perceive sense-data, we should not be obliged to extend this admission to the “normal” cases too. For why on earth should it *not* be the case that, in some few instances, perceiving one sort of thing is exactly like perceiving another? (Austin 1962, p.52)

Austin is pointing out that it doesn’t follow from the fact that E and H are indistinguishable that if in having H you are directly aware of (and only of) a g sense-datum, then in having E you are directly aware of (and only of) a g sense-datum. After all, being directly aware of a g sense-

⁴⁷ We could recast this reasoning as follows. To have a phenomenally conscious experience is to be directly aware of something. This follows from the Confrontation Model and the claim that the objects of direct awareness in

datum could feel from the inside exactly the same as being directly aware of a g avocado.

Though it isn't completely clear, Austin himself seems to have preferred Option 3. (Of course, Austin never would have used the term 'sense-datum.'⁴⁸) Austin's endorsement of this position was one the inspirations for disjunctive naïve realism.

I think the best reason for following the sense-datum theorist in accepting Option 1 is a "uniformity of explanation" principle: if two experiences have the same phenomenal properties, then the account of how one has its phenomenal properties must be the same in relevant respects as the account of how the other has its phenomenal properties.

This is vague, I know, and I will not try to sharpen it here, but the idea behind it is simple. There is something ad hoc about Option 3. If the best explanation of H's G-ness is that it involves your being brutally confronted with an instance of g borne by a sense-datum, then what—other than a desire to preserve a naïve realist account of E—could motivate the decision to explain E's G-ness by saying that it involves brute confrontation with an instance of g borne by a material object? Respect for the uniformity of explanation principle dictates that, once we have established that the best explanation of H's G-ness is that H involves direct awareness of a g sense-datum, we say the same of E. Since we have decided to rule out Option 2—that E involves direct awareness of a g sense-datum *and* a g material object—we are left with Option 1.

For the record, I side with the sense-datum theorist on this question. (I will have more to

perceptual experience are the objects that instantiate the confronted ϕ -properties. But in having H there is no g material object near you. Therefore, you are directly aware of a g immaterial object—a sense-datum.

⁴⁸ Some thoughts on Austin.

Oddly enough, for all of his anti-sense-datum theory bluster, Austin seems to have accepted the first stage of the argument from hallucination. This ought to give pause to those contemporary philosophers who simultaneously denigrate the step and champion Austin's *Sense and Sensibilia*.

One endearing, if slightly irritating, feature of that work is the hilariously disingenuous way in which Austin manages to sneak in commitment to sense-data in cases of hallucination and illusion. On p.32, considering the case of a subject who suffers a hallucinatory water experience while lost in the desert, Austin lambastes sense-datum theorists for saying that the subject is directly aware of a sense-datum. But he does not criticize their reasoning. What is his complaint, then? That we don't need to call what the subject is directly aware of a sense-

say about this in the next chapter.) Therefore, I feel that, once we have come as far as the second stage, the sense-datum theorist is home free: we must say that E, too, involves direct awareness of (and only of) a g sense-datum.⁴⁹

But it doesn't really matter, for the first stage of the argument from hallucination is where the action is. That is for two reasons. First, if that stage does not work, the Option 1 vs. Option 3 issue is moot. Second, the first stage is the one containing the alleged sense-datum fallacy.

The argument from hallucination, then, works as follows. It takes the best explanation of E's and H's phenomenal indistinguishability to be their phenomenal sameness. So H has G. But, since the argument's proponents hold the Confrontation Model, this means that in having H there must be something to bear the ϕ -property instances you confront. That something cannot be material, so it must be immaterial—a sense-datum. End of stage one. By the uniformity of explanation principle, you must be directly aware of a g sense-datum in having E too. Since Option 2 has been ruled out, we can conclude that in having E you are directly aware only of a g sense-datum.

A Sense-Datum Fallacy?

These days, most philosophers believe that the argument from hallucination contains a fallacy in its first stage: the sense-datum fallacy.⁵⁰ Not only do they believe that the argument is fallacious, but they believe that the alleged fallacy is blindingly obvious.

For example, Woozley, in the introduction to his edition of Locke's *Essay*, denies that

datum, for we already have a perfectly good name for it: 'mirage'! Feel free to substitute 'mirage theory' for 'sense-datum theory' throughout this dissertation.

⁴⁹ Smith (2002) contains an extended defense of this "better stop it in the first stage 'cause the second stage works" line of thought.

⁵⁰ I do not know who coined the term "sense-datum fallacy." For some early allegations of it, see Barnes (1965), Chisholm (1957), and Sellars (1956). Ryle (1949) diagnoses the argument from hallucination as suffering from a different fallacy, though I will not discuss Ryle's charge.

Locke held the sense-datum theory, insisting that “It would be hard to understand why anybody should want to rate Locke an important philosopher if his whole theory [the sense-datum or “representative” theory] rests on errors so elementary that a first-year student has no difficulty in spotting them” (1964, p.26-27).⁵¹ And less caustically, but no less insultingly, Harman claims that “the notorious sense datum theory of perception arises through failing to keep these [we will meet them in a moment] elementary points straight” (1990, p.665). The literature on perception and phenomenal consciousness is littered with remarks like these.

I will argue that there is no fallacy in the argument’s first stage. I will also argue that the force of what *does* drive that part of the argument—the Confrontation Model—is felt just as intensely by the sense-datum theory’s greatest critics.

Harman (1990) features a rich discussion of the argument from hallucination and the alleged fallacy. His case against the argument, in its essentials, runs as follows. We can think of the argument’s first stage as moving from (1) it visually appears to a subject S as if there is something of which she is aware to (2) there is something of which S is aware. But this move from “appears as if” to “there is,” Harman thinks, is a perceptual instance of the intentional fallacy involved in moving from (3) Ponce de Leon is searching for something called the Fountain of Youth to (4) there is something called the Fountain of Youth for which Ponce de Leon is searching. Just as it does not follow from the fact that Ponce de Leon is engaged in a doomed search for the Fountain of Youth that there is something for which he is searching, Harman concludes, it does not follow from the fact that it visually appears to S as if there is something of which she is aware that there is something of which she is aware.

What Harman fails to pay proper attention to is that claims to the effect that it visually

⁵¹ I owe this reference to p.8 of Mackie (1976). Chapter 2 of Mackie’s book is an illuminating discussion of Locke’s theory of perception.

appears to a subject as if such and such are *attributions of phenomenally conscious mental events*. This makes them very different from attributions of mental states like beliefs and desires, not to mention states of searching. If it visually appears to a subject as if such and such, then the subject is undergoing a phenomenally conscious visual experience. Sense-datum theorists, in the grip of the Confrontation Model, simply apply their analysis. The Confrontation Model, together with the assumption that the objects of which subjects are directly aware in having their experiences are the objects instantiating the confronted ϕ -properties, demands that if a subject is undergoing a phenomenally conscious visual experience, then there is something of which she is directly aware.

So we can get from Harman's (1) to his (2) without fallacy. Here is the core bit of the stage one reasoning in explicit form:

- P1 If it visually appears to a subject S as if there is something of which she is aware, then S is having a Φ visual experience.
- P2 If S is having a Φ visual experience, then there is a ϕ object of which S is aware.
- P3 It visually appears to S as if there is something of which she is aware.
- C Therefore, there is a ϕ object of which S is aware.

The argument proceeds to conclude that, since the subject could have a Φ experience even if there are no material objects in her vicinity, the ϕ object must be an immaterial sense-datum.⁵²

P1 follows from the fact that visual appearance talk attributes phenomenally conscious

⁵² Lycan accepts P1 through C, but detects a different fallacy in the move from the claim that there is no ϕ *material* object of which the subject is directly aware to the claim that there is a ϕ *immaterial* object of which the subject is aware. He insists that there is an overlooked alternative: that the ϕ object of which S is directly aware is a nonexistent object. See Lycan (2001). (There is also, I suppose, the Russellian alternative that S is directly aware of her own brain. This is implausible for two reasons. First, it's crazy. Second, S's brain need not have the relevant ϕ -property.)

Lycan holds this position because he thinks that phenomenally conscious perceptual experiences are direct awarenesses of *intentional objects*. Thus, in cases of hallucination, the intentional objects of which subjects are directly aware are nonexistent. See later in this section for a criticism of this talk of intentional objects.

experiences to subjects. P2 follows from the Confrontation Model plus the assumption that S is directly aware of the object instantiating the confronted ϕ -property. P3, we can suppose, is a statement of observed fact. The conclusion clearly follows. So far, so valid.

Essentially, then, it follows from the Confrontation Model that whenever it visually appears to a subject as if she is aware of something, the subject is aware of something. In a case of hallucination, there may be no *material* object in the scene before the subject's eyes of which she is aware. Still, there must be *something* of which she is aware, something to bear the property instance she is so plainly brutally confronted with in having her experience. That something must be an immaterial sense-datum.⁵³

There is no sense-datum fallacy, at least not in any of the above reasoning. What there is, though, is the Confrontation Model, something that we ought to reject. So my defense of the sense-datum theory is severely qualified. But philosophers, from the early moderns to the present day, who accept the sense-datum theorists are not fools. They may be wrong, but they aren't making any logical mistakes.

⁵³ Different versions of the sense-datum theory have different accounts of the ϕ -properties.

If the version in question, holds that the ϕ -properties are ordinary perceptible properties like color and shape, properties which both material objects and sense-data can have, it will allow that it follows from the claim that it visually appears to a subject that something is ϕ that the subject is directly aware of something ϕ . Broad was explicit about this:

Whenever I truly judge that x appears to me to have the sensible quality q , what happens is that I am directly aware of a certain object y , which (a) really does have the quality q , and (b) stands in some peculiarly intimate relation, yet to be determined, to x (1923, p.89).

Robinson formulates his own version of this, calling it the "Phenomenal Principle":

(P) If there sensibly appears to a subject to be something which possesses a particular sensible quality then there is something of which the subject is aware which does possess that sensible quality (1994, p.32).

Robinson regards the Phenomenal Principle as the deep principle underlying the arguments from illusion and hallucination. I deny this. The Phenomenal Principle actually rests on the Confrontation Model and the direct awareness assumption. (Chapter 1 of Martin (forthcoming) contains a survey of this aspect of the sense-datum literature. I am indebted throughout this chapter to Martin's work.)

If the version of the sense-datum theory in question insists, as Peacocke did in his (1983), that the ϕ -properties are properties that only sense-data can have, then a slightly modified claim results. Such sense-datum theorists will hold that it follows from the claim that it visually appears to a subject that something is, say, green that the subject is directly aware of something g , where g is a ϕ -property and distinct from green. In a case of a hallucination as of something green, while nothing is green, these philosophers will argue, something is g .

Moreover, Harman himself feels the pull of the Confrontation Model. This is evident in his preference for talk of “intentional objects.” For Harman, phenomenally conscious perceptual experiences are direct awarenesses of intentional objects. If your experience E is veridical, E’s intentional object is the avocado you see. If E is a hallucination, its intentional object is a nonexistent avocado. (It is less clear what he thinks the intentional object is if E is an illusion and the avocado looks r, rather than g, to you.) This is important for two reasons.

For one thing, philosophers should abandon talk of intentional objects of perception, for it only obscures matters. Consider your hallucination H as of an avocado. Harman would say that in having H you are aware of a nonexistent avocado. There are only two ways of understanding this. If we take it literally and regard your experience as a relation between you and something that does not exist, then Harman’s view of experience is obviously false: no one can be related to something that does not exist. If we analyze it out, so that to be aware of a nonexistent avocado is to undergo a mental event with false avocado content, then, while no longer insanely implausible, the view just amounts to the unremarkable claim that experiences have representational content. But then the confusing talk of awareness or perception of intentional objects ought to be dropped in favor of the more tractable idiom of representational content.⁵⁴

I think I know why philosophers go in for this kind of talk, and this leads to the second reason why it is important. First, note that, for the most part, the only philosophers who regularly speak of intentional objects of perception are philosophers attracted to representational

⁵⁴ Anscombe (1965) made talk of perceptual intentional objects famous. I confess that I have never understood her paper. Her “direct objects” of experience have extremely puzzling features.

To return to an earlier footnote, Lycan, too, should abandon his talk of intentional objects of perception. I agree with him that hallucinations are best thought of as relations to false contents, but the nonexistent intentional object idiom only conceals his point.

approaches to phenomenality (either pure or impure representationism). Their diagnosis of H—cleansed of the dark nonexistent object idiom—is that in having it you are in a state with false avocado content. But, naively anyway, just being in a state with false avocado content is not sufficient for having G: your false perceptual belief that there is an avocado before you has false avocado content, but it is not G. By expressing “H has false avocado content” as “H involves direct awareness of a nonexistent avocado,” these philosophers are able to make their representational approach sound more phenomenologically adequate. Why? Because talk of direct awareness of objects evokes the features of experience that motivate the Confrontation Model, especially presentationality. Consider the following passage from a later Harman paper:

Perceptual experience has a certain presentational or representational character, presenting or representing the environment in a certain way. When it looks to you as if you are seeing a ripe tomato, your perceptual experience presents or represents the environment as containing a red and roughly spherical object located at a certain distance and orientation “from here” (1996, p.253).

Here Harman seems half aware of the *prima facie* phenomenological inadequacy of a representational approach to phenomenality, as he switches back and forth from presentation talk (which evokes presentationality and our naïve pre-theoretical understanding of phenomenality) and representation talk (which, to naïve ears, sounds off-key in a discussion of phenomenality, best left for discussion of beliefs and desires).

At all events, I mention this *not* to say either that representational approaches are wrong (we will come to that) or that Harman implicitly holds an application of the Confrontation Model (he does not), but simply to point out that Harman *himself* is moved by the very same pre-theoretical understanding of perception and phenomenality that motivates the Confrontation Model. So not only is the first stage of the argument from hallucination valid, but what makes it go through, the Confrontation Model, isn’t beyond the pale.

5. The Sense-Datum Theory Criticized

Though it does not rest on any sense-datum fallacy, the sense-datum theory is nevertheless highly problematic. Here are some—but by no means all—of the problems that plague it.

First, sense-data raise many metaphysical problems, problems which it is not necessary to rehearse here. Suffice it to say that we would be better off rejecting a theory committed to the existence of things as queer—if it is even clear what sort of thing they are—as sense-data.⁵⁵

Second, the sense-datum theory conflicts with transparency. E is transparent: it has presentationality and locationality, and introspection of it reveals that, in having it, the only properties you are aware of are the properties you are disposed to judge are presented and located—in particular, located in your external, mind-independent environment. g is one of these properties. But the sense-datum theory says that g is instantiated by a sense-datum, not the avocado. Thus, it conflicts with transparency by entailing that g is not instantiated where your experience locates it.

(Moreover, though it is fairly clear how the sense-datum theory has the resources to explain presentationality, it is much less clear how it can explain locationality. Remember that naïve realism explained locationality by pointing to the fact that, according to it, the confronted ϕ -properties are instantiated where subjects judge them to be—in their external, mind-independent environments. But, according to the sense-datum theory, that is precisely where

⁵⁵ On this point, I cannot do better than Bennett:

[Denying the existence of sense-data] also abolishes pseudo-problems. To many readers, writings in the philosophy of perception seem to be addressed partly to (a) genuine and interesting philosophical problems and partly to (b) mildly lunatic conundrums which one would ignore if they did not seem to be forced upon us by theories adduced to solve (a). Among the questions which many would include in class (b), on their own inherent demerits, are these: ‘Can there be an unapprehended sense-datum?’; ‘Are sense-data perceived in the same way, or in the same sense of “perceive”, as are physical things?’; ‘Do apprehended sense-data exist in the same sense of “exist” as do unobserved things?’; ‘Is the visual sense-datum I have just after blinking the very same one that I had just before?’; ‘How do visual sense-data relate to the surfaces of physical objects?’ These ‘problems’ constitute a good part of the class (b) topics in the

they *aren't* instantiated.)

Finally, there is the problem of direct awareness. The sense-datum theory can easily explain:

HALLUCINATION—Hallucinatory experiences—for example, when a subject has a perceptual experience as of an F object in her environment though there is no F object in her environment—are possible.

INDISTINGUISHABILITY—For every veridical experience a subject may have, there is a possible corresponding phenomenally indistinguishable hallucinatory experience.

It does so by maintaining that veridical experiences and their corresponding indistinguishable hallucinations relate their subjects to sense-data instantiating the same ϕ -properties. But this explanation comes at the cost of requiring that subjects are only ever indirectly aware of material objects when they have veridical experiences, and this is in conflict with:

DIRECT AWARENESS—In having veridical experiences, subjects are directly aware of material objects.

Thus, the sense-datum theory is incompatible with direct realism, the view, recall, that subjects can be and usually are directly aware of material objects.

So the sense-datum theory is a form of indirect realism. Though indirect realism was considered the up-to-date, scientifically respectable view of things until relatively recently, these days most philosophers abhor it. Philosophers feel that indirect realism either (1) makes it impossible to have certain kinds of thoughts about material objects or (2) makes it extremely difficult if not impossible to refute the skeptic about the external world.⁵⁶ (1) and (2) are two usual ways of understanding the claim that indirect realism places a problematic “veil of

philosophy of perception which no intelligent adult would put up with if he could see how to disentangle them from (a) the worthwhile topics...” (1971, p.35, footnote omitted).

Bennett leaves out my favorite (b)-type question: do sense-data have backsides?

⁵⁶ On (1), see pp.13-16 in Smith (2002). On (2), see pp.141-147 in Jackson (1977).

perception” between subjects and the external world (see Bennett 1971, p.69). Though I argue elsewhere (see John MS) that indirect realism does not face a veil of perception problem, we ought nevertheless to prefer to the sense-datum theory a theories of phenomenality that do not have indirect realism as a consequence. The intuition that we are directly aware of material objects in perceptual experience is one we should respect.

6. Conclusion

While the sense-datum theory is not the philosophical disaster it is often represented as being, it is a theory we should reject. It commits us to sense-data. It conflicts with the transparency of experience. And it is indirect realist. We can do much better.

Chapter 6: Disjunctive Naïve Realism and the Problem with Brute Confrontation

1. Introduction

There is an application of the Confrontation Model that promises to solve the problem of direct awareness by being compatible with HALLUCINATION, INDISTINGUISHABILITY, and DIRECT AWARENESS. This is disjunctive naïve realism.

I will do two things in this chapter. First, I will present and criticize disjunctive naïve realism. Second, and most importantly, I will argue that no application of the Confrontation Model can be true because the notion of brute confrontation is incoherent.

2. Disjunctive Naïve Realism Introduced and Explained

As an application of the Confrontation Model, disjunctive naïve realism fills in the blanks we found in the statement of the model. It provides an account of what sort of objects instantiate the confronted ϕ -properties and of what the ϕ -properties are.

Disjunctive naïve realism replaces the sense-datum theory's SENSE-DATA and INTRINSICNESS SD with the following two additions to the Confrontation Model:

DISJUNCTION—Either material objects instantiate the ϕ -properties (in the case of veridical experience) or sense-data instantiate them (in the case of nonveridical experience).

INTRINSICNESS DNR—The ϕ -properties are intrinsic properties of their bearers, material objects or sense-data.

This will take a bit of explaining.

Disjunctive naïve realism accepts the first stage of the argument from hallucination, so it grants that if you are undergoing a hallucination H as of an avocado, you are directly aware of a sense-datum. But the theory rejects the second stage of the argument. Instead of concluding

with the sense-datum theorist that in having E as well you are directly aware of (and only of) a g sense-datum, it insists that in having a veridical experience E of an avocado you are directly aware of (and only of) a g material object, an avocado. Following Austin's lead, disjunctive naïve realism chooses Option 3. Disjunctive naïve realism rejects the uniformity of explanation principle that demanded a uniform sense-datum-style treatment of E and H.

Disjunctive naïve realism, like the sense-datum theory, is another attempt to salvage the Confrontation Model in the face of the possibility of phenomenally indistinguishable hallucination. But rather than providing a common account of veridical and nonveridical experience, the disjunctive naïve realist prefers to give the (non-disjunctive) naïve realist explanation of veridical phenomenality while giving the sense-datum theory's explanation of nonveridical phenomenality.

So disjunctive naïve realism says that a phenomenal property is identical with the property of being a mental event whose subject is brutally confronted with an instance of a ϕ -property borne by either a material object or a sense-datum. This view treats perceptual experiences according to the act-object analysis, seeing them as events constituted by relations of direct awareness between subjects and either material objects (in the veridical case) or sense-data (in the nonveridical case). What I said in chapter four about naïve realism is relevant to the theory's account of veridical phenomenality. What I said in chapter five about the sense-datum theory is relevant to the theory's account of nonveridical phenomenality.

Besides Austin, the other founding father of disjunctive naïve realism was Hinton (see Hinton 1973). McDowell (1982, 1986, and 1994), Snowdon (1980-81), Child (1994), Langsam (1995), and Martin (1997 and 2004) all contain defenses of views in the neighborhood of disjunctive naïve realism.

The main motivation for disjunctive naïve realism is that it is an application of the Confrontation Model that manages to solve the problem of direct awareness. It is compatible with all three of HALLUCINATION, INDISTINGUISHABILITY, and DIRECT AWARENESS. It makes room for the possibility of hallucination. It makes room for the possibility of subjectively indistinguishable hallucinations by allowing that veridical experiences and their corresponding subjectively indistinguishable hallucinations have the same phenomenal properties. And it makes room for direct awareness of material objects by explaining veridical experience in naïve realist terms.

Now, if the *only* reason for holding disjunctive naïve realism is that it is compatible with all three of HALLUCINATION, INDISTINGUISHABILITY, and DIRECT AWARENESS, the theory is unmotivated. We will see in the next few chapters that all four applications of the Possession Model solve the problem of direct awareness, and none of them is committed to either the brute confrontation relation or sense-data. However, if the reason for holding disjunctive naïve realism is to preserve the extremely intimate mind-world phenomenal link that naïve realism posits, the theory *is* sufficiently motivated. Unfortunately, that extremely intimate mind-world phenomenal link is a chimera. Though experience can lead us to think that that is how phenomenality is, phenomenality can't be that way—see below.

3. A Complication: There Are Many Disjunctive Theories!

I am not sure how many—if any—of the contemporary philosophers calling themselves “disjunctivists” would accept the theory I have just presented. But this is not a problem. My disjunctive naïve realism is intended as a composite sketch. It captures in only a few strokes several of the philosophical features most important to these philosophers. Though the view I

present looks exactly like none of the extant disjunctivisms, each one of them resembles it a great deal.

Nevertheless, I would like to devote a bit of space to trying to impose some order on the disjunctivist literature. I turn to this project now. I will use ‘disjunctivism’ as a general term for any roughly disjunctive approach to perception and phenomenality. (I will say in a moment what is necessary for an approach to count as disjunctive.)

As Chalmers notes (see his forthcoming), there are at least three disjunctivisms: *experiential disjunctivism*, *content disjunctivism*, and *phenomenal disjunctivism*. I will introduce them, beginning with the last. Let us use ‘V-H pair’ as a term of art for the pair of experiences formed by taking any veridical experience and its corresponding phenomenally indistinguishable hallucination.

Phenomenal disjunctivism says that for any V-H pair the veridical experience will have at least some phenomenal properties that the hallucination lacks (on stronger versions, they will have no phenomenal properties in common). We can think of phenomenal disjunctivism, then, as the view that E and H have different phenomenal properties. Our disjunctive naïve realism is not a form of phenomenal disjunctivism, since it allows that E and H can have all of the same phenomenal properties.

Content disjunctivism says that for any V-H pair the veridical experience will have at least some representational content that the hallucination lacks. We can think of content disjunctivism, then, as the view that E and H have different contents. Our disjunctive naïve realism is not a form of content disjunctivism. If it is compatible with naïve realism that experiences have content, then it is compatible with disjunctive naïve realism that E and H have all of their content in common. But, as I remarked in chapter four, that experiences have content

may be incompatible with naïve realism. Either way, our theory is not content disjunctivist.⁵⁷

Experiential disjunctivism is the view that for any V-H pair the subject is in a type of mental state in having the veridical experience that she is not in in having the hallucination. We can think of experiential disjunctivism, then, as the view that you are in a mental state in having E that you are not in in having H. There are two ways to develop experiential disjunctivism.

The experiential disjunctivist could argue for (1) a naïve realist treatment of veridical experience and the special mental state you enter into only in veridical perception and (2) a specific non-naïve realist treatment of nonveridical experience and the distinct mental state you enter into in nonveridical perception. Our disjunctive naïve realism is this sort of experiential disjunctivism. It says that in having E you are in a mental state that is your being directly aware of the avocado; it says that in having H you are in a mental state that is your being directly aware of a sense-datum. The identity conditions of each of these mental states are object-dependent: they make reference to the objects they are directed on.

Alternatively, the experiential disjunctivist could argue for (1) a naïve realist treatment of veridical experience and the special mental state you enter into only in veridical perception but (2) refuse to say anything—at least anything *philosophical* (there could be a great many scientific things to say)—about what takes place in nonveridical perception. Though our disjunctive naïve realism is not this sort of view, many forms of disjunctivism are, such as Martin's (see his 1997). (Williamson, though not a disjunctivist himself, claims that disjunctivists should embrace this second sort of experiential disjunctivism. See pp.47-48 of his 2000.) They say that in having E you are in a mental state that is your being directly aware of

⁵⁷ Perhaps it could turn out that disjunctive naïve realism is trivially content disjunctivist. Suppose naïve realism is incompatible with allowing that experiences have content. But, as we saw in chapter 5, the sense-datum theory is compatible with content. Thus, it could turn out that according to disjunctive naïve realism veridical experiences lack content and nonveridical experiences have it—content disjunctivism on the cheap.

the avocado; they say that in having H you are either in a different mental state whose nature is left undescribed or in no mental state at all—you only seem to be in a mental state (the mental state of being directly aware of the avocado).

Experiential disjunctivism is probably the most widely held disjunctive approach. Indeed, it is the approach that inspired the label ‘disjunctivism.’ Philosophers long thought, disjunctivists argue, that the best explanation of E’s and H’s phenomenal indistinguishability was that, in having each, you are in the same mental state. This mental state is a “common element” or, as it is put in McDowell (1982), a “highest common factor” between E and H. Call the traditional view “conjunctivism.” Conjunctivism says that it visually appears to a subject S as if there is a material object before her if and only if she is in a mental state M of a certain type, one which she can be in even if she is hallucinating. What determines whether its visually appearing to S this way counts as seeing a material object or merely hallucinating one is the presence or absence of the appropriate sort of causal link between the material object and M. Experiential disjunctivism, on the other hand, denies that there is a highest common factor, a mental state serving as a common element in E and H. It says that it visually appears to S as if there is a material object before her if and only if *either* (1) she sees a material object (and is thus in a mental state she can be in only if her experience is veridical) *or* (2) she hallucinates (and is thus either in a different mental state or in no mental state at all).^{58,59}

⁵⁸ Actually, not even this formulation of the difference between conjunctivism and disjunctivism is adequate. The problem is that most disjunctivists wish to reserve special treatment only for fully veridical experiences. They throw hallucinations and illusions on the “either a different mental state or no mental state at all” heap. But the formulation I give above would allow illusions to receive the special treatment reserved for fully veridical experiences.

So here’s a better formulation: according to experiential disjunctivism, S has a visual experience as of something’s being F if and only if either (1) S sees that something is F (and is in a mental state she can be in only if her experience is fully veridical) or (2) it merely seems to S as if something is F, in which case her experience is either an illusion or a hallucination (and she is in either an illusion/hallucination mental state or no mental state at all). This expresses the desired contrast with conjunctivism and groups illusions with hallucinations.

I draw these distinctions to point out just how unclear the issue of disjunctivism is. There are many ways of being a disjunctivist, and it is not always clear which one any particular self-professed disjunctivist endorses.

Here is a reason, one of many, why that lack of clarity matters. Sometimes disjunctivism is presented as phenomenal disjunctivism. Phenomenal disjunctivism sees the flaw in the argument from hallucination in the claim that E and H have the same phenomenal properties. But as we saw in the previous chapter, it really does seem as if that claim furnishes the best explanation of their phenomenal indistinguishability. I submit that disjunctivism is best understood as experiential disjunctivism, for one can hold it and grant that E and H have the same phenomenal properties. What is fundamental for disjunctivism, I think, is not anything to do with phenomenal properties but rather the way in which we explain how experiences have their phenomenal properties. As we have seen, experiential disjunctivists reject the uniformity of explanation principle that demands we give the same account to E and H. Disjunctive naïve realism, a form of experiential disjunctivism, rejects the principle, insisting that a material object instantiates *g* in E and a sense-datum instantiates it in H.

4. Disjunctive Naïve Realism Criticized

This section presents some problems for disjunctive naïve realism.

To begin with, disjunctive naïve realism inherits the flaws of both naïve realism (minus

⁵⁹ What are the logical relations between these three forms of disjunctivism? Experiential disjunctivism entails neither phenomenal nor content disjunctivism. One could hold that in having E you are in a mental state you could not be in if there were no avocado before you, while allowing that E and H have all of their phenomenal properties and content in common. (This might be a strange view, but it isn't incoherent.) Phenomenal property and content disjunctivism are logically independent of each other if phenomenality and intentionality are also logically independent of each other. I believe that they are, but we will not be in a position to assess this claim until the end of the dissertation. Finally, do either phenomenal property or content disjunctivism entail experiential disjunctivism? I tend to think that they do, but I am unresolved.

those tied to naïve realism's failure to deal with hallucination) and the sense-datum theory (minus those tied to the sense-datum theory's failure to secure DIRECT AWARENESS). We do not need to recapitulate them. See the relevant sections of chapters four and five.

One worry facing disjunctive naïve realism is that it is unclear how the theory can deal with illusion, what occurs when a subject succeeds in perceiving an object but perceives it as having a property it actually lacks. Suppose the avocado sits on a tabletop in the corner of an Ames Room. As a result, the avocado looks enormous. You succeed in seeing the avocado, but it visually appears to you to be many times larger than its actual size.

The disjunctive naïve realist, as far as I can see, can say one of either two things.

First, he could say that, in having your experience of the avocado, you are not directly aware of the avocado. Rather, you are directly aware of a sense-datum bearing the ϕ -property instantiations necessary for you to have an experience as of a giant avocado. This is the same as the sense-datum theory's explanation. But surely even in illusions you remain directly aware of the avocado. This reply to the worry cannot allow that.

Second, the disjunctive naïve realist could say that you are directly aware of two objects, the avocado and a sense-datum. But now we face some difficult questions. Though the avocado visually appears to you to be much larger than it really is, for the most part it appears the same to you—it looks the same, only a lot bigger. Presumably only the sense-datum bears the ϕ -property instantiations responsible for the larger-than-life size appearance. But what of the other, veridically perceived properties? Is it the avocado, the sense-datum, or both that instantiates them? And how are the avocado and the sense-datum related? Why don't you seem to be aware of two things? Is your dual direct awareness of them somehow phenomenally fused into apparent direct awareness of single enormous-looking object? It would be best if we didn't have

to answer these questions.

Another worry about disjunctive naïve realism is whether it has any scope for application. The theory says that veridical experiences involve brute confrontation with properties instantiated by material objects, while non-veridical experiences, illusions and hallucinations alike, involve brute confrontation with properties instantiated by sense-data. But if—as is certainly possible and likely to be actual—most perceptual experiences are shot-through with illusion, then there are no veridical—at least no *perfectly* veridical—experiences to which it would be appropriate to give the disjunctivist’s special naïve realist account. Disjunctive naïve realism, then, would end up in practice furnishing precisely the same account of perception and phenomenality as the sense-datum theory it was supposed to supplant.

(This possibility is particularly worrisome for those forms of experiential disjunctivism that refuse to give any account at all of nonveridical experience. If there are no perfectly veridical experiences, it could turn out that the disjunctivist is left with nothing philosophical to say about perception. Perhaps McDowell, drawn to such a “quietism” in the moral realm (see his 1994), would not mind this. But surely there is something philosophical to say about perceptual experience and phenomenality.)

5. The Trouble with Brute Confrontation

But disjunctive naïve realism’s greatest flaw is one it shares with naïve realism and the sense-datum theory—indeed, with *any* possible application of the Confrontation Model: its commitment to brute confrontation.⁶⁰ This section will argue that the brute confrontation

⁶⁰ You might have been wondering about applications of the Confrontation Model other than the three we have discussed. We will meet a fourth application in a later chapter. Hint: it says that the objects instantiating the ϕ -properties are our experiences themselves.

relation is incoherent. If that's right, then we should abandon the Confrontation Model—even if we find ways to patch its applications' other flaws.

Brute confrontation is brute because it is unanalyzable, in particular, unanalyzable in representational terms: for any ϕ -property instance, a subject's being brutally confronted with it is *not* to be analyzed as the subject's being in a mental state with the content that ... ϕ We also saw that brute confrontation is supposed to be a cognitive relation: in standing in it to a property instance, the mind somehow grasps the property in a way that puts the subject in a position to have propositional knowledge. Just by seeing the avocado and being brutally confronted with an instance of g , you are thereby in a position to think thoughts and form beliefs concerning g . You might not wish to. You might even be prevented by lack of attention or damage to your brain or visual system. But you will be in a position to do these things. By contrast, standing in the three-feet-away-from relation to the avocado doesn't suffice to put you in a position to do these things. That is because the three-feet-away-from relation is not cognitive.

So the brute confrontation relation is supposed to be both nonrepresentational and cognitive. The problem, though, is that no nonrepresentational relation can be cognitive: all cognitive relations—all relations between minds and property instances which are such that standing in them enables the subject whose mind it is to have the ability to form propositional knowledge—are representational. It follows that there can be no such relation as brute confrontation. It would have to be representational and nonrepresentational.

It is important to see that brute confrontation must be both nonrepresentational and cognitive if it is to do the work the Confrontation Model needs it to do.

Suppose the brute confrontation were representational. Then a subject's being brutally

confronted with an instance of ϕ would be the subject's being in a state with the representational content that ... ϕ But to the philosopher attracted to the Confrontation Model, such a representational relation would be useless for explaining phenomenality. Take E. If your being brutally confronted by an instance of g in having E is just your experience's representing that g is instantiated, then, worries our philosopher, what distinguishes E from the occurrent thought you might form about the avocado, a thought with precisely the same content? We come once more to our naïve suspicion of representational approaches to phenomenality—more on this in the next few chapters.

So if any relation is to do what the philosopher who wishes to give a Confrontation Model-style account of phenomenality wants, an account which respects our naïve anti-representational prejudice, the relation must not be a representational relation.

Why must brute confrontation be cognitive? Because no matter what else you think of the relation that obtains between your mind and the instance of g when you have E, you simply must think that, simply in virtue of having E, you are thereby in a position to form beliefs of the form that ...g.... Now, to say this is not to say that simply in virtue of having E you ought to be able to form beliefs of the form that ...g... *which you would be able to express using language*. After all, surely a non-colorblind animal with the cognitive sophistication of, say, a dog is able to form beliefs of the form that ...g...—if it helps, to divide the space of possibilities into the g and not-g possibilities—just in virtue of seeing the avocado.⁶¹

⁶¹ Two things.

First, again, the claim is that having E puts you in a position to know, not that it is sufficient for you to know. Other cases illustrate this. You are confronted with so much rich detail in having E that you would have to focus your visual attention in various ways to form certain thoughts and beliefs. But E puts you in a position to do this. Standing next to the avocado with your eyes closed doesn't.

Second, suppose I ask you to look carefully at your hand and focus on the salient surface property of it that we would be inclined to call its color. Then I say, "surely the relation you stand in to this property instance is cognitive." Note something special about this demonstration: you are *introspecting* your experience of your hand in performing it. Perhaps, then, the relation R you stand in to your hand's color while introspecting your experience of

But at this point an objection arises: why are only representational relations cognitive? I can imagine someone saying: “I agree that the philosopher who endorses the Confrontation Model had better say that brute confrontation is nonrepresentational and that brute confrontation is cognitive. But why can’t there be nonrepresentational yet, for all that, cognitive relations? After all, doesn’t standing in the eyes-open-looking-right-at-the-avocado-good daylight-actively-focusing-on-the-avocado’s-salient-surface-property relation—call it ‘@’—to an instance of g put you in a position to have propositional knowledge?”

Reply: Yes, standing in @ to g does suffice to put you in a position to have propositional knowledge. But that is only because standing in @ to g entails standing in a genuinely cognitive relation to g. Think of it this way. There is a lot of clearly non-cognitive stuff involved in standing in @ to g. The eyes open part, for example, isn’t sufficient for you to stand in a cognitive relation to g: you could have your eyes open but be in a coma or asleep. Likewise for just about every other part of @. What is crucial, though, is that standing in @ to g involves a focusing-on-the-avocado’s-salient-surface-property part, and that necessarily involves having a visual experience as of something g. And it simply is not plausible that that experience puts you into a cognitive relation to g in any other way than by its involving your representing something as being g. Without this, though, @ would not be cognitive: it would be no more cognitive than the standing-next-to relation. So it can’t be standing-next-to all the way down.

I am leaning on intuition here, I know. But this intuition is backed up by something more than representationalist prejudice. At bottom, what is necessary for a relation between a subject

your hand is different from the relation R^* you stand in to your hand’s color while not introspecting your hand. If so, and if our only access is to R^* , how can we be confident that R is cognitive? Maybe R^* is, but what about R ? I do not know how to answer this question.

This raises a general worry. Many of the claims we make in the philosophy of mind are based on introspection of our experiences. But what is our justification for thinking that the way our experiences are when we introspect them is the way they are when we are not introspecting them? I would like to investigate this issue in future work.

and something else to be cognitive is that the subject's standing in that relation to the something suffices to confer correctness conditions on relevant mental events and states of the subject. In short, a mental relation is cognitive only if a subject's standing in it to something is sufficient for one or more of the subject's mental events or states to be correct or incorrect. But that is just for the mental relation to be representational.⁶²

6. Conclusion

Now we are in a position to appreciate what is wrong with the Confrontation Model. First, each of its applications is implausible on its own terms. I dealt with these difficulties in this and the previous two chapters. Second, and more importantly, the Confrontation Model's brute confrontation relation is incoherent. I argued for this claim in the foregoing section.

In the next chapter, we turn to the Possession Model, a conception of phenomenality which does away with the brute confrontation relation altogether.

⁶² I suspect that my argument is veering towards the celebrated attack on the "Myth of the Given" in Sellars (1956). I hope to explore in the future the connections between my argument against brute confrontation and Sellars's argument against the Given.

I also suspect that my argument against brute confrontation applies equally to the theory of appearing.

Chapter 7: The Possession Model

1. Introduction

This chapter introduces the Possession Model and says a bit about its history. After that, it discusses four ways in which it has been applied. Finally, it assuages a worry about representationalist accounts of perceptual experience.

2. The Possession Model

The Possession Model denies that all there is for a mental event to have a phenomenal property is for the event's subject to brutally confront a property instance. It rejects commitment to brute confrontation relations and posits an entirely different conception of phenomenality. The Possession Model holds the following:

POSSESSION MODEL—What it is for a mental event to have a phenomenal property Φ is for the event to have a non-confrontation-based property.

Thus, the Possession Model denies that G is constituted by your brute confrontation with an instance of g . It simply identifies phenomenal properties with special non-confrontation-based properties. So, according to the Possession Model, your experience E 's phenomenal property G is identical with a special non-confrontation-based property. A mental event's having that property is necessary and sufficient for it to have G .

The Possession Model furnishes a mere schema, one which can be filled in in different ways by specifying different candidate non-brute confrontation-based properties for identification with the phenomenal properties. Different applications of the Possession Model make different recommendations as to which properties we ought to select. I will discuss four

applications of the Possession Model: pure representationism, impure representationism, impure phenomenism, and pure phenomenism.

Note that, unlike the Confrontation Model's schema, the Possession Model's schema makes no explicit reference to the *subjects* of phenomenally conscious mental events. The reason for this, I think, is that the Possession Model does not view phenomenality as constituted by relations between subjects and property instances (or anything else). It sees phenomenality simply as a matter of mental events having certain special properties.

3. From the Confrontation Model to the Possession Model

Because of the presentationality of phenomenally conscious perceptual experience, we all feel, whether we recognize it or not, the pull of the Confrontation Model. But the Possession Model lacks the sort of strong phenomenological motivation that bolsters its rival (though, as we will see below, it isn't without *some* phenomenological support). Indeed, the motivation for the Possession Model is largely negative: a desire to avoid, not just particular applications of the Confrontation Model, but the entire way of thinking about phenomenality the Confrontation Model enshrines.

Dissatisfaction with the sense-datum theory led some philosophers to question the act-object analysis of experience (see Ducasse 1951 and Chisholm 1957). In place of that analysis, philosophers gradually adopted what has since come to be called the "adverbial" analysis of experience. The act-object analysis views experiences as mental events constituted by direct awareness relations between subjects and objects of perception, either sense-data (in the case of the sense-datum theory) or material objects (in the case of naïve realism and disjunctive naïve realism). The adverbial analysis, on the other hand, views experiences as mental events marked,

not by their being constituted by direct awareness relations (they are not, according to the analysis), but by their possession of certain sorts of properties.

The difference between the two analyses is illustrated in the different accounts they would give of the claim that in having your experience E of the avocado you seem to see something g. The act-object analysis would say that you seem to see something g if and only if you are directly aware of a g object. But the adverbial analysis would deny this account, substituting it with: you seem to see something g if and only if you visually sense g-ly. Here nothing is said about objects of awareness. What *is* said is that you visually sense in the way you typically do when you successfully visually sense—that is, see—g things.⁶³

But what way is that? Adverbial accounts express the thought that perceptual experiences and other phenomenally conscious mental events have phenomenal properties. Thus, to visually sense g-ly is to visually sense in the way you typically do when you see g things, and that is to have a veridical G visual experience, where G is a non-brute confrontation-based property of experience. The rejection of the act-object analysis leads to a rejection of the Confrontation Model of phenomenality that generated it. By rejecting the act-object analysis, the adverbialist is able to see the possibility that phenomenality might consist, not in brute confrontation with property instances, but in possession of a certain sort of (non-brute confrontation-based property).

⁶³ Note how this echoes the “type-neutral” descriptions of phenomenally conscious mental events in Smart (1959). Smart’s paper even contains a brief attack on Confrontation-style explanations of phenomenality (though he didn’t put it in those terms).

Speaking of Smart, the first chapter of Martin (forthcoming) features an insightful discussion of how the mid-twentieth century rise of physicalist reductionist explanations of phenomenal consciousness also played a major role in bringing about the shift towards Possession-style explanations of phenomenality.

The turn away from act-object analyses and towards adverbialism, then, reflects a turn away from Confrontation-style phenomenal property explanations towards various Possession-style explanations. Note two features of the Possession Model.

First, the Confrontation Model, as we saw in chapter four, regards the phenomenal properties of experiences, and thus experiences themselves, as only *derivatively* phenomenal. What is really phenomenal—really “lights on!”—in the world are ϕ -property instances, and experiences (and other phenomenally conscious mental events) become “phenomenal” only when their subjects are brutally confronted with those instances. The Possession Model, by contrast, sees phenomenally conscious mental events *themselves* as the fundamental “phenomenal things” in the world.

Second, we will see that applications of the Possession Model will have no problem solving the problem of direct awareness. Thus, each is easily compatible with direct realism. Confrontation-style explanations of phenomenality have trouble with this because of the way they make standing in the direct awareness relation to something (the bearer of the confronted ϕ -property instances) necessary for phenomenality. This leads to pressure to find something in hallucination that subjects can be directly aware of, and that, in turn, leads to pressure to make whatever it is the subject is aware of in hallucination available in veridical perception—no more direct realism. The Possession Model changes all of that. Its applications will view the question of direct awareness—the question of what subjects are directly aware of in having their experiences—as the simple question: what objects and properties do subjects perceive in having their experiences?

Viewing early- and mid-twentieth century philosophy of perception through the confrontation vs. possession prism also helps explain another curious feature of that period: the

almost overnight switch on the part of most philosophers of perception from thinking of perceptual experiences as being in the first instance relations to “objects of perception” (usually sense-data) to thinking of them as being relations to “contents” (propositions).⁶⁴ This is an enormous topic that others have written a great deal about (see, for example, the papers collected in Crane 1992), so let me say just this. As we will see below, most applications of the Possession Model identify phenomenal properties with either pure or impure representational properties. To say that a mental event has a representational property, though, is to say, among other things, that it is a relation to a representational content, a proposition. Since the Possession Model does not think phenomenality essentially requires direct awareness, we have a view that treats perceptual experiences as relations to propositions.

Philosophers made the transition from “object of perception” thinking to “content of perception” thinking in fits and starts. But the transition from the Confrontation Model to representationalist applications of the Possession Model was eased by talk of perceptual experiences as direct awarenesses of “intentional objects” (see Anscombe 1965). As I argued in chapter 5, this is a problematic way of speaking of perceptual experience. It eases the confrontation-to-possession transition by couching the core representationalist idea, that phenomenality is intentional, in terms of awareness of objects—something with a distinctly “confrontational” air. This is problematic because it obscures the true nature of representationalist applications of the Possession Model. We will come back to this in the chapters on pure and impure representationism.⁶⁵

⁶⁴ Of course, philosophers spoke for centuries of the “contents” of experience, but they meant by ‘content’ particular objects such as material objects or sense-data, not propositions. If you like, simply revise the above paragraph so that it asks why there was such a sudden transition from regarding the contents of experience as particular objects to regarding them as propositions.

⁶⁵ Other philosophers with representationalist sympathies, philosophers such as Armstrong (see Armstrong 1968) and Pitcher (see Pitcher 1970), analyzed perceptual experiences as dispositions to form beliefs. Such views strike

We can view these representationalist applications of the Possession Model as attempting to correct a phenomenological inadequacy in the Confrontation Model. Something that kept coming up with each application of the model was the failure to adequately explain locationality. We can *sort of* see how someone might have thought that brute confrontation explains presentationality. But locationality was always trouble. Pure and impure representationists, as we will see in the next two chapters, insist that locationality is explained by the fact that experiences have content according to which various properties are instantiated in subjects' environments. Why, on this view, are you disposed to judge that E locates g? Because E represents a content to the effect that g is instantiated in your external, mind-independent environment.⁶⁶

Today most of the battles over phenomenality are fights over which application of the Possession Model to adopt. Indeed, with one exception (the theory of appearing, which I have described in previous chapters as a contemporary application of the Confrontation Model), every theory of phenomenality taken seriously these days is an application of the Possession Model. This reveals something odd about the current philosophical state of play. In one sense, philosophers have not come to a broad consensus regarding the best way to answer what I called in the introduction the "structure question." There is no explicit agreement on how we ought to think about the structure of experience and phenomenality. Yet, in another sense, philosophers *have* reached a consensus, for, though they would not put it this way, they all subscribe to the Possession Model.

me as ways of ignoring—if not outright denying—phenomenality rather than explaining it. See pp.45-47 of Smith (2002) for discussion of the curious tendency of some philosophers to deny phenomenality.

⁶⁶ I suspect that perceptual phenomenology generates a tension: presentationality motivates the Confrontation Model, while locationality motivates the Possession Model. I hope to explore this idea in future work.

I doubt that a physicalist reductionist answer can be given for what I called in the introduction the “reduction question” if we accept the Confrontation Model. The reason is simple: brute confrontation is supposed to be primitive and unanalyzable. This looks like dualist non-reductionism, for there can be no explanation of brute confrontation in terms derived solely from the physical sciences. If, say, the sense-datum theory is true, then completed science will have to mention unanalyzed brute confrontation.

There are at least two respects in which philosophers with representationalist sympathies would profit from a look at the confrontation-possession dialectic even though that dialectic has a de facto, if undeclared, winner. These philosophers often seem to misunderstand the current range of live theoretical options. We have seen how they sometimes try to sell their positions by using rhetoric best suited to a defense of an application of the Confrontation Model like naïve realism. But they also tend erroneously to view non-representationalist applications of the Possession Model as holdout versions of the sense-datum theory. We will see that they are not.

4. The Four Applications

At last, it is time to introduce the four applications of the Possession Model we will discuss: pure representationism, impure representationism, impure phenomenism, and pure phenomenism.

Each of these four theories of phenomenality is an answer to the question “Can phenomenality be analyzed in terms of intentionality?” Pure representationism and impure representationism say “yes”; impure phenomenism says “sorta”; pure phenomenism says “no.”

We can think of these four theories as lying along a continuum. At one extreme of the continuum are theories that seek to explain phenomenality in entirely representational terms. At

the opposite extreme are theories that seek to explain phenomenality in entirely nonrepresentational terms.

At the extreme representationalist end, we find pure representationism, which accounts for phenomenal properties in wholly representational terms. It identifies phenomenal properties like G with pure representational properties like the property of representing the content that something is g. It says: G = the property of representing the content that something is g.

Slightly less representationalist is impure representationism, which accounts for phenomenal properties in mixed representational and nonrepresentational terms (heavy on the representational terms). It identifies phenomenal properties like G with impure representational properties like the property of representing the content that something is g in a visual manner. It says: G = the property of representing the content that something is g in a visual manner.

Even less representationalist is impure phenomenism, which accounts for phenomenal properties in mixed representational and nonrepresentational terms (heavy on the non-representational terms). It identifies phenomenal properties like G with complex properties composed of nonrepresentational properties of experiences together with some pure representational property or other. It allows the nonrepresentational and pure representational properties to vary independently of each other. It says: G = the property of representing some content or other and having a nonrepresentational property. Impure phenomenism often refers to its favored nonrepresentational properties as “qualia.”

Finally, we find pure phenomenism at the extreme nonrepresentationalist end of the continuum. It accounts for phenomenal properties in wholly nonrepresentational terms, identifying phenomenal properties like G with nonrepresentational properties of experience. It

says: G = a nonrepresentational property. Pure phenomenism also often refers to its favored nonrepresentational properties as “qualia.”

It will be easier to understand our four applications of the Possession Model if we compare each with a thesis that has come to be called intentionalism.

Intentionalism is the claim that the phenomenal properties of a mental event supervene on the event’s representational properties; that is, there can be no phenomenal difference between two mental events unless there is a representational difference between them.⁶⁷

We can draw a useful distinction between two forms of intentionalism. On the one hand, there is *inter-modal* intentionalism, which says that, for any two mental events regardless of their mental type (occurrent thought, visual experience, itch, act of imagining a Rothko), if the mental events have the same representational properties, then they have the same phenomenal properties. *Intra-modal* intentionalism, on the other hand, says that, for any two mental events of the same mental type, if the mental states have the same representational properties, then they have the same phenomenal properties.⁶⁸

Think of things this way. Inter-modal intentionalism says that, *across* mental state types, phenomenal differences supervene on representational differences. So, according to it, the phenomenal differences between E and T, a tactile experience of the avocado, supervene on representational differences between them. Intra-modal intentionalism says that, *within* any particular mental state type, phenomenal differences supervene on representational differences. So, according to it, the phenomenal difference between a G visual experience and an R visual experience supervenes on a representational difference between them. Inter-modal

⁶⁷ Byrne (2001) defends intentionalism.

⁶⁸ The labels ‘inter-modal’ and ‘intra-modal’ and this way of using them to distinguish among varieties of intentionalism are Byrne’s.

intentionalism entails intra-modal intentionalism, but not vice versa.

Notice that if spectrum inversion without misrepresentation is possible, then both forms of intentionalism are false. Consider the following scenario. Your experience E represents the avocado as being g and is a G experience. Jones's experience E* also represents the avocado as being g, but it is an R experience. Here we have a case of two visual experiences, two mental events of the same mental type, with the same content but different phenomenal properties. If this case is possible, intra-modal intentionalism is false. And if intra-modal intentionalism is false, inter-modal intentionalism is false too.⁶⁹

Pure representationism entails inter-modal intentionalism (and, thus, intra-modal intentionalism). But not vice versa, as pure representationism is an identity—not just a supervenience—claim. Since it identifies phenomenal properties with representational properties, it makes phenomenal differences supervene on representational differences *across* mental state types. This means that the possibility of spectrum inversion without misrepresentation is enough to rule out pure representationism (unless special steps are taken—more on these in the next two chapters).

Matters are a bit more complicated when it comes to impure representationism. A major question facing impure representationists is whether to allow that mental events of different types can share content. Can your experience E and your occurrent thought T share content, say, that something is avocado-shaped? I will have much more to say about this issue in chapter 9. For now, let us simply note this much. Impure representationism plus the assumption that mental events of different mental types cannot share content, entails intermodal intentionalism *if we assume that, for any two mental events we consider, both have phenomenality*. (If one were

⁶⁹ See Shoemaker (1982), Harman (1990), and Block (1990) for classic discussions of spectrum inversion and its importance for intentionalism.

to lack it, the phenomenal difference between them would have to supervene on a difference not in content but in specific manner of representation: in, in other words, one of their nonrepresentational properties.) Otherwise, if we allow (as seems plausible) that mental events of different types can share content, then impure representationism entails intra-modal intentionalism. Since it identifies phenomenal properties with impure representational properties (properties of representing properties in specific mental manners), it makes phenomenal differences supervene on representational differences *within* particular mental state types. Again, this means that the possibility of spectrum inversion without misrepresentation is enough to rule out impure representationism (unless special steps are taken—more on these in the next two chapters).

Impure phenomenism entails neither inter-modal nor intra-modal intentionalism. (A theory's entailing either inter- or intra-modal intentionalism is necessary for it warrant the label 'representationism.')

Though it makes a mental event's possession of (pure) representational properties necessary for its possession of phenomenal properties, it does not tie those phenomenal properties to the representational properties by having the latter determine the former. The possibility of spectrum inversion without misrepresentation is no problem for impure phenomenism. Impure phenomenism allows an experience to have G even if it has r or b or y—and not g—content. As long as it has some content or other, if the experience has the requisite nonrepresentational property, it has G.

Pure phenomenism also entails neither inter-modal nor intra-modal intentionalism. It not only allows phenomenal properties and representational properties to float free of one another, as does impure phenomenism, but it denies that a mental event must possess representational properties to possess phenomenal properties. The possibility of spectrum inversion without

misrepresentation is no problem for pure phenomenism either.

Thus, three of our four applications of the Possession Model hold that phenomenality can be explained in terms of intentionality, while one denies this. Note two important things.

First, to say that three of the four theories hold that phenomenality can be explained in terms of intentionality is not to say that, according to these theories, phenomenality can be *reduced* to intentionality. If, for example, it turns out that intentionality ca not be understood without recourse to phenomenality-involving notions, then the intentional analysis will not suffice for an intentional reduction. Below, when I examine each of these theories in turn, I will say much more about the explanation-reduction question. Also note that pure phenomenism is not non-reductive simply because it does not analyze phenomenality in terms of intentionality. For all I have said, “qualia,” the nonrepresentational properties it identifies with phenomenal properties, could be functional or physical properties of subjects’ mental events.

Second, the theories I will discuss disagree over whether—and if so, to what extent—phenomenality can be explained in terms of intentionality. But there is another school of thought gaining adherents these days, one which says that rather than looking to ground phenomenality in intentionality we ought to ground intentionality in phenomenality.

A full discussion of this view lies well beyond the scope of this dissertation, so I will limit myself to the following provisional remarks. (Warning! The foremost proponent of this position is Charles Siewert. Unfortunately, I have only recently begun attempting to come to grips with his important work (see his 1998). Thus, what I say about this view will be directed at a very brief summary of it in Chalmers 2004a. I am well aware that this is not an ideal way of approaching the position. I will correct this in future work.)

To quote Chalmers’s recent summary of this position, “there is at least a distinctive and

crucial sort of intentional content that accrues in virtue of the phenomenal character of mental states. These theorists do not typically offer reductive analysis [sic] of intentionality in terms of consciousness, but they typically hold that consciousness has a certain priority in the constitution of intentionality” (2004a, p.1, footnote omitted). In addition to Siewert and Chalmers, McGinn and Searle defend this view (see McGinn’s 1991 and Searle’s 1983).

Chalmers continues with the following:

[The following thesis] seems plausible: that any two phenomenally identical states will share some aspect of their representational content. As Colin McGinn has put it, there is a strong intuition that something about representational content is *internal* to phenomenology. That is, given a specific phenomenology, it seems that if a mental state has this phenomenology, it must also have a certain specific representational content. If this is right, then phenomenal properties entail certain pure representational properties (ibid, pp.3-4, emphasis in original).

Then he goes on to summarize Siewert’s argument for this claim:

This line of thought has been developed into an argument at greatest length by Charles Siewert (1998). I will not recapitulate Siewert’s extensive arguments here, but the basic idea is that visual experiences are *assessible for accuracy*, in virtue of their phenomenal character. For example, when I have a visual experience as of something X-shaped in front of me, this experience may be either accurate or inaccurate, depending on what is really in front of me. Further, it seems that any visual experience with the same phenomenal character would be assessible for accuracy in the same sort of way: roughly, it would be correct if and only if there is an object with an appropriate shape in front of the subject. If this is right, then a phenomenal property (having an experience with the phenomenal character as of seeing something X-shaped) entails a pure representational property (roughly, representing that there is an object with a certain shape in the world) (ibid, p.4, emphasis in original).

What should we think about this argument?

It seems to go as follows: if a perceptual experience has a phenomenal property, then, simply in virtue of having that property, the experience is assessible for accuracy; if two perceptual experiences have the same phenomenal property, then, simply in virtue of having that property, they are assessible for accuracy in the same way; but if that is right, then any experience with a phenomenal property has a pure representational property (say, the pure representational property of representing the content that would be true if and only if the experience’s accuracy condition, determined solely by its having a particular phenomenal

property, is satisfied).

I think we ought to be suspicious of this. Why be so sure that experiences are assessible for accuracy solely in virtue of their phenomenal properties? We must distinguish the *appearance* of assessibility for accuracy from really being assessible for accuracy. There is something about phenomenally conscious perceptual experiences (and other phenomenally conscious mental events as well, but we can ignore them) that makes them *seem* as if they are assessible for accuracy solely in virtue of their phenomenal properties. This something is just what I have been calling the locational aspect of their phenomenality. This is indeed an essential feature of any phenomenally conscious mental event. But it doesn't follow from this that phenomenally conscious experiences *are* assessible for accuracy in virtue of their phenomenal properties. If something is assessible for accuracy, it is such that it has a correctness condition. For something to have a correctness condition, though, it is necessary for that thing to have had a certain sort of history, one in which the thing enjoyed the right kinds of causal connections to things in the world. But as long as it is possible for something to have phenomenal properties without having had this sort of causal history, it is possible for something to seem to be assessible for accuracy without being so.

To sum up, this argument has two problems. First, it mistakes the appearance of assessibility for accuracy for the real thing. That we can see this is another benefit of distinguishing the presentational and locational aspects of phenomenality. Second, it does not follow from an experience's seeming to be assessible for accuracy that it is. Suppose phenomenal internalism—on which phenomenal properties supervene on subjects' intrinsic properties—and content externalism—on which representational properties do *not* supervene on subjects' intrinsic properties—are both true. Then Swampman, say, could have a G

experience—thus having an experience that *seems*, in virtue of its locationality, as if it is assessible for accuracy—even though he arguably enjoys no states with any content—thus being in no states that *are* assessible for accuracy. Locationality is such that having it makes a mental event *feel* contentful. It does not follow from this that it is.

Two final comments.

First, I am leaning very hard on (1) the intuition that phenomenal properties are narrow and (2) the intuition that representational properties are broad. I'll defend both in chapter 9. For the record, though, the only way I can possibly see the appearance of contentfulness entailing contentfulness is if either (a) both phenomenal internalism and content *internalism* are true or (b) both phenomenal *externalism* and content externalism are true.

Second, I believe that locationality lies behind much of the suspicion of externalist theories of intentionality. The idea is this. Most philosophers are phenomenal internalists even if they are content externalists. (Putnam and Burge, for example, both seem to have been phenomenal internalists.) But a perceptual experience's having phenomenality entails its having locationality. Locationality, as I have noted, can be described as the feeling of contentfulness. Thus, philosophers skeptical of content externalism can point to locationality-exhibiting phenomenally conscious perceptual experiences and say the following: "Look at these mental events. They have content and that content is fixed by subjects' intrinsic properties. So content externalism must be false. Either there's no such thing as broad content or, at worst, there is narrow content." I claim that these philosophers would be mistaking locationality, an aspect of phenomenality that may be empty of content, for intentionality. I will not argue for this claim; this is just speculation about possible deep motives for resistance to content externalism.

In addition to not discussing the “phenomenality grounds intentionality” view, I will not discuss so-called “higher-order thought” and “higher-order perception” representationalist theories of phenomenality. See Byrne (2004) for more on these important views. They also lie too far beyond the scope of this dissertation.

5. A Worry About Representational Accounts of Experience

All four applications of the Possession Model we will discuss, even pure phenomenism, hold that phenomenally conscious perceptual experiences (at least usually) have representational content. In this section, I will discuss a worry that is sometimes raised about accounts that regard experiences as relations to propositions.

Our naïve—naïve because naïve realist at heart—intuition is that phenomenally conscious perceptual experiences literally present properties to us in all their phenomenal splendor. I have tried to capture this intuition by speaking of brute confrontation with property instances. The naïve intuition about paradigmatically representational states (e.g. beliefs), however, is that they serve, not to present properties to us, but literally to *represent* them when they are not before our minds for presentation in all their glory. Your experience E, so the naïve thought goes, puts you in direct contact with g, while your perceptual belief B about the avocado you see, had with your eyes closed, is only a record of the direct contact with g you once enjoyed in experience. The naïve thought can allow that B relates you to g—just not in the intimate way in which E relates you to it.

A metaphor might help cash out the naïve intuition. The naïve view sees perceptual experience as a sort of open space in which subjects’ minds come into direct contact with properties. Think of this open space as rather like a large room in which a party is being held, a

party at which subjects' minds and a motley assortment of perceptible properties are the guests. Experiences, then, are like parties at which our minds march right up to properties, introduce themselves, and shake hands. Beliefs, on the other hand, are records—representations—of the properties met. They are like business cards, things which we can take away from the party so that, even when the properties we met are not there to shake hands with, we can remain connected to them. Here, shaking hands with a property is the analog of being related to a property in a phenomenally conscious way.⁷⁰

So, on the naïve view, it sounds wrong to say that E represents g, for it regards E as a kind of mental *encounter* with g, not a mental event or state (such as a belief) whose role is to serve as a *memorial* of an encounter with g.⁷¹

The worry should now be clear: perceptual experiences just can't be representational. Thinking of them as representational would just be a category mistake. They are not the kinds of mental events or states that can be representational.

My response to this worry should be obvious. This naïve contrast between perceptual experience and paradigmatically representational mental states like belief rests on commitment to the Confrontation Model. Once we abandon that model, though, we can see our way clear to

⁷⁰ Perhaps this metaphor sheds some light on the view, held by philosophers from Aristotle to the early modern period, that the mind must in a quite literal sense have to touch—shake hands with!—something in order to perceive it. This was part of the idea behind ancient extromission theories of perception. It also motivated the rhetoric of presence to mind that played such an important role in early modern discussions of perception. See Arnauld (1990).

⁷¹ One needn't be naïve to feel the force of this naïve contrast. Alston, sophisticated as they come, writes:
First your eyes are shut, and you are *thinking* about the scene in front of you, say your front yard. You *remember* the trees in the yard. You *wonder* whether there are squirrels and robins out there at the moment. You *hypothesize* that your neighbor across the street is working in his garden. . . . Then you open your eyes and take a look. Your cognitive condition is radically transformed. Whereas before you were just *thinking* about, *wondering* about, *remembering* the trees, the squirrels, the houses, and so on, these items (or some of them) are now *directly presented* to your awareness. They are *given* to your consciousness. They are *present* to you, whereas before you were merely dealing with propositions about them (1999, p.72, emphasis Alston's).

This comes from a paper defending the theory of appearing. And in an earlier paper, Alston argued, citing among other things this very contrast, that perceptual experiences are not “conceptual,” by which he seems to have meant intentional. See Alston (1998).

the thought that perceptual experiences, initial appearances notwithstanding, aren't something fundamentally different from representational events and states.

However, the naïve worry should give pause to pure and impure representationists. These days, the phenomenal pendulum has swung so far from the Confrontation Model side to the Possession Model side that philosophers with representationalist sympathies sometimes write as if having content were sufficient for phenomenality. Indeed, pure representationism actually accepts this. (Impure representationists don't, but, as we will see, they often write as if they do.) The naïve contrast teaches us an important lesson: though experiences may be representational, intentionality is not sufficient for phenomenality. In the rush away from Confrontation-style explanations of phenomenality, this lesson was sometimes lost.⁷²

6. Conclusion

Now that I have introduced the Possession Model, it is time to turn to its four applications. While the Possession Model provides a schema explaining what it is for a mental event to have a phenomenal property, its applications provide detailed theories of phenomenal properties.

The plan for the rest of this dissertation is as follows. I will devote a chapter to pure representationism, a chapter to impure representationism, and a chapter to pure and impure phenomenism. I will argue by elimination for pure phenomenism.

⁷² Here is McGinn criticizing Dretske's representationalist approach to phenomenality:

The upshot is that there is a large gap at the centre of Dretske's account of consciousness. He simply has no articulated account of it, once one looks closely; indeed, he never really even squares up to the questions that need to be raised. He simply *asserts* for the most part that the nature of conscious experience consists in its having naturally representational content, without giving us any convincing account of what makes such states conscious to begin with. . . . Dretske seems so intent on defending a representational theory of qualia that he virtually ignores the prior question of what makes an internal state a conscious state to begin with. . . (1997, p.531, emphasis McGinn's).

McGinn is accusing Dretske of having failed to pay proper respect to the lesson taught by the naïve intuition.

Chapter 8: Pure Representationism

1. Introduction

In this chapter, I will introduce and argue against pure representationism. We will see that pure representationism suffers from difficulties curiously like those suffered by naïve realism.

I will proceed as follows. I will begin with a section outlining the motivation for a representationalist approach—pure or impure—to phenomenality and the divergent ways in which pure and impure representationism act on that motivation. Then I will present a series of criticisms of pure representationism.

2. Why Be a Representationist?

Philosophers attracted to representationalist approaches to phenomenality usually offer two distinct reasons for holding their positions, one broadly theoretical and the other phenomenological.

First, the broadly theoretical consideration. There are many mind-body problems. In addition to the problem of phenomenality (and of subjectivity), there is the problem of intentionality.⁷³ The problem of intentionality is the problem of explaining how certain mental events and states can represent or be “about” objects, properties, and states of affairs.

So the problem of intentionality can be thought of as the problem of explaining how certain mental events and states can have representational properties, just as the problem of phenomenality can be thought of as the problem of explaining how certain mental states can

⁷³ See Stich and Warfield (1994) for papers on the problem of intentionality and Block, Flanagan, and Güzeldere (1997) for papers on the problem of phenomenal consciousness.

have phenomenal properties.⁷⁴

This way of describing the problems reveals the broadly theoretical motivation for pure and impure representationism's identification of phenomenal properties with representational properties. Philosophers tend to regard the problem of phenomenality as the most difficult of the mind-body problems, so difficult that some philosophers despair of our ever solving it (see McGinn 1989). But most philosophers believe that the problem of intentionality is much easier, and some are confident that we are nearing a solution to it (see, among many others, Dretske 1981, Fodor 1987, Lewis 1994, and Stalnaker 1984). If phenomenal properties are representational properties, though, the problem of phenomenality is just a special case of the problem of intentionality. Arguably, a reduction of the problem of phenomenality to a special case of the problem of intentionality would render the former problem no more difficult than the latter.⁷⁵ Thus, if true, a representationalist approach would tract the seemingly intractable problem of phenomenality and provide a physicalist reductionist answer to the reduction question. Or so pure and impure representationists hope.⁷⁶

Now for the phenomenological motivation. The transparency of experience is supposed to argue against pure and impure phenomenism, leaving pure and impure representationism. In chapter 10, I will argue that the argument from transparency fails.

There are other representationalist arguments. Chief among these is the argument in Byrne (2001) for intentionalism. But even if his argument is successful, it is only an argument

⁷⁴ There are two sides to the problem of intentionality (just as there are two sides, explaining presentationality and explaining locationality, to the problem of phenomenality): that of explaining what it is for a mental event or state to have content (pure representational properties) and that of explaining what it is for a subject to bear a particular attitude towards a content (impure representational properties).

⁷⁵ Arguably, because one might wonder, as McGinn does (see his 1991), whether this wouldn't simply show that intentionality is much more problematic than we previously thought.

⁷⁶ Impure phenomenism, though it makes having content necessary for having phenomenality, is not sufficient for the representationalist's reductive aim. That is because it does not accept intentionalism, the supervenience of the phenomenal on the intentional.

for intentionalism, which, merely claiming that the phenomenal *supervenes* on the representational, won't suffice for pure representationism.

Note that pure representationism is compatible with direct realism. In having a veridical experience, you are directly aware of the object(s) you perceive, and your experience has only true content. In having an illusion, you are directly aware of the object(s) you perceive, but your experience misrepresents it (them). In having a hallucination, you are directly aware of no object(s), though, because of your experience's false content, you seem to be. So pure representationism can solve the problem of direct awareness.

3. What is Pure Representationism?

Recall that pure representationism identifies phenomenal properties like G with pure representational properties like the property of representing the content that something is g. So, according to pure representationism, a mental event has G if and only if it represents that something is g. Thau (2002) is the only explicit defense of pure representationism I know of, though we will see that impure representationists sometimes write as if they hold the pure form of representationism.

By identifying phenomenal properties with pure representational properties, pure representationism has to account for phenomenality entirely in terms of the contents represented by phenomenally conscious mental events, rather than in terms of both content and attitude (specific manner of representation). As we will see, given this, the only way for pure representationism to have a chance at being true, is for it to appeal to very special contents, contents involving a very special class of properties.

Pure representationism says that perceptual experiences are propositional attitudes.

4. Against Pure Representationism

This section raises a series of difficulties for pure representationism.

The Problem of Common Content

Take pure representationism's identification of G with the pure representational property of representing that something is g. If g is just green, the avocado's color, pure representationism is obviously false. If representing green were sufficient for having G, your belief that the avocado is green, your occurrent thought that the avocado is more green than the avocado you saw yesterday, the interpreted sentence 'something is green', and smoke signals indicating green would all have G. They don't, so G is not the property of representing that something is green. Note that this conclusion holds not just for the colors and phenomenal properties corresponding to color experience. It holds for all of the perceptible properties (color, shape, distance, etc.) and the phenomenal properties corresponding to them.

Let us call this problem, after Egan (2004), the "problem of common content": G, say, cannot be identical with a pure representational property involving representation of a feature—such as color—that a mental event or state *without G* can also represent.⁷⁷ The problem of common content makes plain that, if one wants to hold pure representationism, one must identify properties like G with pure representational properties involving representation of properties that are distinct from the colors, shapes, distances, etc.⁷⁸

This is exactly what Thau does (see Thau 2002). According to Thau, G = the property of

⁷⁷ The problem of common content is the analogue of the modality problem that plagued naïve realism. See chapter 4.

⁷⁸ Just to be clear, the reason is that in virtue of its being *pure* representationism, the theory must explain phenomenal properties solely in terms of properties represented. Thus, it can't do what impure representationism can do: let nonrepresentational properties (the properties that fix the specific manner of representation) determine the phenomenal modality within which differences in pure representational properties will fix phenomenal differences.

representing that something is *g*, where *g* is a property “distinct from but corresponding to” green (p.35). Let us call this property “Thau-green.” Call the special class of perceptible properties it belongs to the “Thau-properties.”⁷⁹

Thau-green is a very special property. It has two important features:

- (1) Only the human visual system can represent it.
- (2) Representing it is necessary and sufficient for having *G*.

Neither of these things is true of green. Your occurrent thought *T* about the avocado represents green even though it is not a mental event generated by the human visual system. And *T* represents green without having *G*.

(2) explains why *E* has *G*. *E* represents Thau-green, so *E* has *G*. (1) and (2) explain why *T* does not have *G*. *T* is not a mental event generated by the human visual system, so *T* cannot represent Thau-green. But if *T* cannot represent Thau-green, *T* can't have *G*. In short: (2) follows from Thau's form of pure representationism; (1) is what avoids the problem of common content.

So one can hold pure representationism only if one holds that the relevant represented properties are Thau-properties. For the rest of this section, then, when I speak of pure representationism, I mean the view that phenomenal properties are identical with pure representational properties involving the representation of Thau-properties. Should we believe pure representationism when understood this way? I will argue that we should not.

⁷⁹ Thau also thinks that visual experiences do not represent the colors, shapes, distances, etc. of things, but only the Thau-properties—Thau-colors, Thau-shapes, Thau-distances, etc.—of things. Thus, he denies that colors, shapes, distances, etc. are perceptible properties. However, he doesn't deny that things have colors, shapes, distances, etc. For the moment anyway, this striking part of his position won't concern us. What matters to us is his identification of phenomenal properties with pure representational properties involving representation of Thau-properties.

However, I will point out that there is a tension between this aspect of Thau's theory and the naïve realist-sounding motivations he sometimes adduces for it.

What You Can See You Can Think About!

The first objection to Thau's pure representationism isn't so much an objection as it is a call for more explanation: just what are Thau-properties? We know that (1) and (2) above hold of them, but not much more. Moreover, Thau's case for his theory must be ironclad for us to believe that such properties exist. Ontological scrupulousness together with the intuition that the salient surface property of the avocado that causes you to sort it with zucchini, not with pomegranates, militates against belief in them.

The really serious objection to pure representationism is that it violates the following principle:

REPRESENTATION PRINCIPLE—Any property that can be (purely) represented by a perceptual experience can be represented by a thought, belief, desire, wish, fear, etc.

This principle says that if your experience E can represent the avocado as being Thau-green, then you can have a thought or belief or desire or wish or fear that represents the avocado as being Thau-green: what you can see you can think about. If Thau is right about the Thau-properties, though, the Representation Principle is false, for Thau-properties are supposed to be perceptually representable but not, say, doxastically representable.

So Thau must reject the Representation Principle. Why is that such a big deal?

For one thing, rejecting the Representation Principle is *prima facie* implausible. Suppose that after leaving the avocado behind you begin thinking about its salient surface property g. (You were so impressed by the avocado's rich, dark g-ness that you can't help but think about it.) Denying the Representation Principle means denying that the property you are so enthusiastically thinking of—the property we are inclined to call the avocado's color—is the

property you *saw*, the property the avocado *looked* to you to have. That is bizarre.

It's bizarre because we tend to think that perceptual representation of perceptible properties is *cognitively promiscuous*. As Evans (1980) remarked, we are “informational systems” (p.122). We take in information through our senses. That information is traded between our various cognitive systems. Some of the information we take in is sent back out through speech or other communicative behavior. We model this information flow by talking of shared representational contents. If Thau is right, perceptual representation of what we thought were colors, shapes, distances, etc. is taken in through the senses, but it can never be passed on to other cognitive systems. Since Thau presumably thinks that material objects actually have Thau-properties (though he isn't always clear), this means that information concerning the objects in our environments—information concerning their Thau-colors, Thau-shapes, Thau-distances from other Thau-propriety objects, etc.—is taken in through the senses only to...only to stop right there. That information can never become the content of any thoughts, beliefs, or desires.

Given this, it is not clear how Thau's theory can even allow the possibility of subjects gaining information about their environments through their senses.⁸⁰ That is very bad.

Nonconceptual Content to the Rescue?

Philosophers who defend the claim that perceptual experiences have “nonconceptual” content (e.g. Evans 1980, Peacocke 1992, Tye 1995 and 2000) might come to Thau's defense.

⁸⁰ Two points.

First, *don't* say that Thau's theory can allow perceptual information gathering as follows: perceptual experiences take in their special contents, then nonperceptual systems infer from these contents the distal instantiation of colors, shapes, distances, etc. This will not work for a simple reason. In order for the nonperceptual states to perform the inferences, they would have to be able to share content with perceptual experiences. For if they could not do this, they could not register, as first premises in their inferences, propositions like the proposition that perceptual experiences with such-and-such content have occurred. In short, if Thau's theory is true, we are informationally trapped behind our senses.

Perhaps the pure representational properties Thau needs are pure representational properties involving nonconceptual content. Let us explore this strategy.

There are two plausible interpretations of the claim that experiences have non-conceptual content: the state view and the content view.⁸¹

On the state view, there is only one sort of content: propositions—whatever they turn out to be (for example, sets of worlds, Fregean Thoughts, Russellian singular propositions). A content—a perfectly ordinary content—will count as nonconceptual if and only if a subject can be in a mental state with that content even if the subject is unable to deploy in thought the concepts that a suitable interpreter would use in characterizing the subject's attitudes.

On the content view, nonconceptual contents are a different sort of content—that is, proposition—than conceptual contents. In particular, a content is nonconceptual if and only if it is not a Fregean Thought, where Thoughts are somehow compositionally built up out of Fregean senses or “concepts.”⁸² (So nonconceptual contents will be sets of worlds or Russellian propositions.)⁸³

Now, what Thau needs are a class of contents—propositions—concerning some properties—the Thau properties—to which subjects can be related only in certain mental events and states. Why? Because he thinks it is sufficient for a mental event or state to have a phenomenal property that it be related to one of these contents. So if either the state or content view is to help him, it must entail that in having E you are related to a proposition concerning

Second, it is implausible that there are properties which are such that they only permit themselves to be represented, so to speak, by a specific type of state of only a specific kind of thing, humans.

⁸¹ See Byrne (2003) and Heck (2000) for helpful discussion. Stalnaker (1998) argues that, in at least one clear and important sense, *all content*—perceptual and nonperceptual—is nonconceptual.

⁸² Frege didn't call his senses “concepts,” of course. (He meant by ‘concept’ something closer to what we mean by ‘property.’) But the most vocal friends of the claim that experiences have conceptual content (see, for example, McDowell 1994) are Fregeans who think that concepts just are Fregean senses.

⁸³ The terms ‘conceptual content’ and ‘nonconceptual content’ originate in Evans (1980).

Thau-green to which you can be related only in having visual experiences.

The state view is obviously of no help to Thau. If the state view is right, then in having visual experiences subjects are related to propositions they can be related to even if they cannot deploy in thought the concepts a suitably placed interpreter would use to characterize their visual attitudes. But it does not follow that these propositions are ones they can be related to only in having visual experiences. To see this, consider the person we can call “Know It All.” Know It All is extremely conceptually sophisticated: for every property there is, he has a concept of it, and he is capable of deploying that concept in thought. Know It All, therefore, can be related in thought to every one of the propositions he can be related to in perceptual experience. (Thus, Know It All’s perceptual beliefs about the avocado ought, if Thau were to supplement his theory with the state view, to be G. But they aren’t.) So the appeal to the state view would not help Thau reject the Representation Principle.

Does the content view fare any better? If it is right, then, in having E, you are related to a non-Fregean content: a Russellian singular proposition, say—call it “C.”

First of all, it does not follow from this that you are not *also* related to C in thinking your occurrent non-G thought T about the avocado. What is crucial to defenders of nonconceptual content is not that merely doxastic mental events and states cannot be relations to nonconceptual contents, but that perceptual experiences (and, as we will see in a moment, subpersonal representational states) cannot be relations to *conceptual* contents. So the content view does not provide aid and comfort to Thau’s pure representationist denial of the Representation Principle.

But, second, even if we just *stipulate* that according to the content view only perceptual experiences can be relations to nonconceptual contents, problems ensue.

To begin with, believers in nonconceptual content need it to be the case that unconscious

subpersonal states of our various cognitive systems can share nonconceptual content with perceptual experiences. To deny this would be to deny that we are informational systems, landing us right back where we started with Thau's implausible rejection of the Representation Principle. That principle really does seem non-negotiable. So much the worse for any theory, nonconceptualist or otherwise, that entails its falsity.

Moreover, it does not follow from the claim that only perceptual experiences can be relations to nonconceptual contents that perceptual experiences and merely doxastic mental events and states can't have some content in common. What philosophers are doing when they say that only perceptual experiences can be relations to nonconceptual contents conceived as, say, Russellian singular propositions is attempting to model the way in which perceptual experiences are informationally richer than thoughts and beliefs.⁸⁴ We might say that E relates you to a singular proposition in order to capture E's richness. *But that is a claim about the type of entity—singular proposition or set of worlds or Fregean Thought—to which you are related in E, not a claim ruling out that E and your occurrent thought T both represent the avocado as being g.* Talk about whether any given mental event is related to a Russellian singular proposition or a set of worlds or a Fregean Thought is, at bottom, philosophical shop talk. It is talk about which sort of entity we should appeal to in order to express in our theories certain otherwise difficult to express facts about mental life. So saying that E is and T is not related to a nonconceptual content, a Russellian singular proposition, is perfectly compatible with saying that they both represent many of the same properties. (It is just that E, due to its richness, represents so many *more* properties than T does.) The content view, then, cannot help Thau. He needs it to be the case that E and T share no g (that is, Thau-green) content in common to prevent T from

⁸⁴ Though they do not *have* to say that only perceptual experiences are relations to these contents to do this. They just need to say that perceptual experiences don't have conceptual content—see above.

being G. But it does not follow from the content view, even if we strengthen it in the way suggested above, that E and T do not have some g content in common. All that follows is that E is a relation to a type of entity to which T cannot be related.

So the appeal to nonconceptual content cannot help.

Does Thau Really Reject the Representation Principle?

Rejecting the Representation Principle is a highly implausible move to make. But Thau would reply to this by insisting that he *does* hold the Representation Principle. Thau grants that as long as you are either seeing or imagining something Thau-green at time t, you can have thoughts, beliefs, desires, fears, etc. at t that represent Thau-green. So *while you are looking at the avocado* (or imagining it) you can think the thought that something has *that property* (namely, the salient surface property the avocado looks to you to have). Of course, you cannot think this thought after you have stopped seeing (or imagining) the avocado.

It is not clear why Thau makes this move. To respond to the charge that he violates the Representation Principle, Thau must show that, on his position, subjects can at least sometimes represent Thau-properties in thought, belief, and desire. To do this, though, he needs to produce a case in which a thought represents Thau-green (and, by pure representationism, has G). Described as neutrally as possible, though, the case he actually produces is one in which (1) a subject both undergoes a visual experience E of the avocado (or an act of imagining of it I) while entertaining an occurrent thought T and (2) at least one of the two mental events—E or T—has G. Question: what is the most plausible interpretation of Thau's case? It is clearly one in which the experience—not the thought—has G and the thought—itself without G—merely *accompanies* the experience. But if this is right, Thau has not produced a case that shows his

view to be compatible with the Representation Principle.⁸⁵

Even if he had come up with such a case we would still want an explanation of why nonperceptual (and non-imaginative) mental events can represent Thau-properties only when they accompany perceptual (or imaginative) events. Moreover, the spirit of the Representation Principle clearly entails that we ought to be able to have thoughts that represent Thau-properties pretty much anytime we wish, not only when we are having accompanying experiences. So Thau's position is still in conflict with the principle.

I think what leads Thau astray here is his reliance on talk of "acquaintance." He writes:

What makes a perception a perception—and what makes it a conscious episode—is the properties it represents. . . . But, in order for some object to be represented as having one of the intrinsic surface properties we mistake for the colors, such acquaintance [visual or imaginative acquaintance] is required; *what it is for one of these properties to be represented to you is for you to be visually or imaginatively acquainted with it* (2002, p.233, emphasis Thau's).

Thau seems to confuse the "visual acquaintance" with Thau-green that the experience provides with the "visual acquaintance" his theory says is necessary for representation of Thau-green. But, to remain consistent with his view, Thau would have to say in the case we are considering that there are *two* instances of visual acquaintance: the one furnished by the experience and the other furnished by the accompanying thought. He can't say that the thought *does* get to represent Thau-green by pointing to the visual acquaintance that exists in the case unless he can show that the visual acquaintance he points to is the *thought's* visual acquaintance. But the most natural reading of the case makes the visual acquaintance the *experience's*.

⁸⁵ There is another problem with Thau's line. He grants that you can have a Thau-green-representing thought while either visually experiencing Thau-green or imagining it. Note, though, that there is an undeniable phenomenal difference between seeing Thau-green and imagining it. But if Thau's pure representationism is true, this phenomenal difference shouldn't exist. He needs to say, then, that the experience and the act of imagining represent *different* properties. But now it looks as if Thau must deny that you can even imagine Thau-green. For, if you could, your act of imagining would be phenomenally the same as your experience of Thau-green. So Thau has to believe that there are Thau-properties that can only be represented in perceptual experience *and* that there are distinct properties—super Thau-properties!—that can only be represented in imagination. Not only, on his view, can you not think about (or imagine) the properties you see, but you can't think about the properties you imagine either.

Here is the worry put another way. Thau holds that what it is to represent Thau-green is for you to be visually acquainted with Thau-green and to be visually acquainted with Thau-green is to have a G visual experience. So, by this, if your thought T represents Thau-green, then you are visually acquainted with it, i.e., you have a G visual experience. But having a thought cannot be sufficient for having an experience! (You can decide to think that p, but you can't decide to be visually acquainted with anything.) What we know about the case he describes is that there is some visual acquaintance with Thau-green. But this is *visual* acquaintance! So it is most plausible that the mental state in virtue of which there is some visual acquaintance is E, not T. Again, T just accompanies E.)⁸⁶

From Pure to Impure Representationism

But even if we set aside the problems that the Representation Principle poses for Thau's position, his view faces yet another problem.

Recall that an experience's having a phenomenal property entails its being such as to dispose its subject to judge that it makes a property present to mind. Phenomenality entails presentationality. Pure representationism says that phenomenal properties are identical with certain pure representational properties. As we have seen, these properties will have to be pure representational properties involving the representation of Thau-properties. Thus, pure representationists must believe that a mental event's having a pure representational property involving the representation of a Thau-property entails its having a phenomenal property and

⁸⁶ Two points.

First, note how close to naïve realism Thau comes in the quoted passage.

Second, Thau offers no justification for the claim that "what makes a perception a perception—and what makes it a conscious episode—is the properties it represents..." (ibid). Indeed, he seems throughout much of his book to assume the truth of pure representationism. (This is especially clear in his response on pp.222-226 to an important objection to his theory posed by Byrne.) I ought to point out that Thau (2002) isn't really a defense of

hence presentationality, but that a mental event's having a pure representational property involving the representation of a non-Thau-property does *not* entail its having a phenomenal property and hence presentationality. In other words, pure representationists are committed to the existence of a division in the class of pure representational properties between those whose possession entails presentationality and those whose possession doesn't entail presentationality.

What explains the existence of this division? That is, what explains the difference between the pure representational properties identified with phenomenal properties and those not so identified? Why does a mental event's having a property of the first sort entail its being such as to make its subject spontaneously inclined to speak of presence to mind while a mental event's having a property of the second sort does not?

Indeed, for all Thau has told us, the existence of the division is a mystery. Presumably, Thau thinks that 'represents' means the same in both 'This mental event represents Thau-green' and 'This mental event represents green.' And Thau-green and green are, as properties, on a metaphysical par. Both belong, for example, to the avocado's surface. (Though, again, it is not always clear how Thau thinks of the Thau-properties.) Why, then, does the representation of one but not the other suffice for presentationality?

There seem to be only two options for Thau. Either it is a brute fact that representing Thau-green suffices for presentationality while representing green does not. Or there is something about the visual manner in which Thau-green is represented that explains why its being represented—*visually* represented—suffices for presentationality.

The first option is implausible. Why should the representation of one surface property of the avocado be so radically unlike the representation of another surface property of the avocado?

pure representationism. He has other fish to fry, and pure representationism is invoked at various places to make other arguments work.

Surely there is a non-brute answer to this question.

The second option is plausible, but it constitutes an abandonment of pure representationism. Remember, pure representationism entails that a mental event has G if and only if it represents that something is g. But the second option scraps this biconditional for another: a mental event has G if and only if it *visually* represents that something is g—represents that something is g *in a visual manner*. The second option is just impure representationism. Thau, then, should give up his implausible pure representationism, with its problematic commitment to Thau-properties, and take up impure representationism—along with all of the other philosophers sympathetic to representational approaches to phenomenality.⁸⁷

Let us call this final problem the “problem of presentation”: pure representationism must either hold that it’s a brute fact that pure representation of Thau-properties entails presentationality while pure representation of colors doesn’t or become impure representationism.

5. Conclusion

Pure representationism is very implausible. Why have I spent so much time on it? For two reasons.

First, because we learn some important lessons along the way. For example, we learn that granting experiences “nonconceptual content” isn’t a viable pure representationist strategy.

Second, because I want to show exactly how untenable pure representationism is. Philosophers with representationalist sympathies often write as if they hold pure representationism; their rhetoric, at any rate, certainly supports this interpretation. For example,

⁸⁷ This problem is an analog to the specifying property problem that plagued naïve realism. See chapter 4.

Dretske, an impure representationist, insists that “[a]ll mental facts are representational facts” (1995a, p.xiii). He seems to have facts about pure representational properties in mind. And Lycan goes so far as to refer to his theory, again, a form of impure representationism, as “the hegemony of representation” (1996, p.11).⁸⁸ Perhaps these philosophers feel as if pure representationism is the view they would hold if only they had a rainy day to work all the bugs out. But they shouldn’t feel this way if they do. There can be no completely representationalist theory of phenomenality. If we want to be representationists, it is impure representationism or nothing.

⁸⁸ I want to stress that both Dretske and Lycan are proponents of impure representationism, not pure representationism. This makes their use of such rhetoric, which sounds decidedly pro-pure representationism, quite puzzling.

Chapter 9: Impure Representationism

1. Introduction

In this chapter, I will introduce and argue against impure representationism.

2. What is Impure Representationism?

Impure representationism aims to reap the theoretical benefits promised by a representationalist account of phenomenality while avoiding pure representationism's implausible consequences. It does this by identifying phenomenal properties like *G* with impure representational properties like the property of *visually* representing the content that something is *g*. So, according to impure representationism, a mental event has *G* if and only if it represents that something is *g* in a visual manner. Thus, impure representationism appeals to two independently necessary and jointly sufficient conditions for phenomenality: possession of the requisite pure representational property (the property in virtue of which a mental event represents its content) and possession of the requisite nonrepresentational property (the property in virtue of which a mental event represents its content in the specific manner—takes the attitude towards its content—in which it does).

Impure representationism says that experiences are special sorts of propositional attitudes. Like pure representationism, it is able to offer a direct realist solution to the problem of direct awareness.

3. How Do Impure Representational Properties Help?

Let me say a bit about why, exactly, impure representational properties help.

Let us return briefly to pure representationism's problems. It says that G = the property of representing the content that something is g . But if g = green, then thoughts, beliefs, desires, smoke signals, etc. with the content that something is g should be G too, but they aren't. This was the problem of common content. It forces the pure representationist to identify properties like g with special properties, the Thau-properties. But this move had two problems of its own. First, it rendered pure representationism incompatible with the unimpeachable Representation Principle. Second, it left it unclear how there could be any non-mysterious account of why representation of Thau-properties, but not other properties, was sufficient for presentationality.

The impure representationist recognizes what is wrong with pure representationism: its commitment to the claim that representation of a property—*any property*: a color like green or a Thau-color like Thau-green—is sufficient for phenomenality. The impure representationist agrees with the pure representationist that representing that something is g is *necessary* for having G . But it cannot be sufficient.

In place of pure representationism, the impure representationist will say the following. What is sufficient for having G is the property of *phenomenally* representing that something is g . To represent phenomenally that something is g is to represent that something is g *in a phenomenal manner*. This is why the impure representationist identifies phenomenal properties with impure representational properties (here, the property of representing a content in a phenomenal manner). To speak of a phenomenal manner of representation is to make explanatory appeal to more than just content. It is to make explanatory appeal to content *and* the particular attitude towards that content. My belief that p and my desire that p both have the pure representational property of representing that p . But only the belief has the impure representational property of representing that p in the belief manner. Why? Because only it is a

mental state in which I bear the belief attitude towards p. Likewise, your experience E of the avocado and your perceptual belief B about the avocado both have the pure representational property of representing that something is g. But, according to the impure representationist, only E has the impure representational property of representing that something is green in the phenomenal manner. Why? Because only it is a mental state in which you bear the phenomenal representation attitude towards p.

Benefits follow immediately from the impure representationist's position. For one thing, by appealing to a special phenomenal manner of representation the impure representationist can allow both that (1) E and B represent that something is g and that (2) g = green. Phenomenal differences are no longer determined solely by differences in content, but by differences in content and attitude towards content. So E and B can have the content that something is g (where g is green) in common without B's having to have G. This solves the problem of common content and allows impure representationism and the Representation Principle to coexist. The problem of presentation is also easily dealt with, for the impure representationist is no longer committed to a curious division in the class of pure representational properties. (See the last chapter for more on the curious division.)

Other benefits accrue. In virtue of its appeal to impure representational properties we can see impure representationism as providing the resources to account for phenomenality's two aspects, presentationality and locationality. By making the possession of representational content necessary for phenomenality, it has the resources to account for locationality. Why are you disposed to judge that E locates g in your external, mind-independent environment? Because, says impure representationism, E represents g as instantiated by the avocado. And by making special perceptual phenomenal manners of representation necessary for phenomenality,

it has the resources to account for presentationality. Why are you disposed to judge that E makes g present to your mind? Because, says impure representationism, the nonrepresentational properties your experience possesses in virtue of which it visually (phenomenally) represents g determine that g seems, well, presented in having E.⁸⁹

4. Two Forms of Impure Representationism

At this point, you are probably thinking something like this: “Wait a minute! You began by characterizing impure representationism as the view that phenomenal properties are identical with impure representational properties like the property of visually representing that something is g. But now, in explaining impure representationism, you say that, according to it, phenomenal properties are identical with impure representational properties like the property of *phenomenally* representing that something is g. Why the difference?” The answer to this question points to a distinction between two forms of impure representationism.

I have been considering an *experience-thought* problem of common content: the problem that arises when we say both that (1) experiences and “merely doxastic” events and states can share content and that (2) having content is sufficient for phenomenality. It is enough to deal with this problem of common content by appeal to phenomenal manners of representation.

But there is also a *sense modality-sense modality* problem of common content. Suppose we identify phenomenal properties with impure representational properties of phenomenally representing contents and allow that experiences in different sense modalities can have content in common. Take V, your visual experience of the avocado’s shape, and T, your tactile experience

⁸⁹ Remember: the point here isn’t that impure representationism explains *everything*. It still leaves us, arguably, with the famous explanatory gap. The point is that, unlike pure representationism, impure representationism has the theoretical sources to cover all of the phenomena that need explaining.

of the avocado's shape. V and T plausibly represent the avocado as having the same shape property; so V and T both represent that something is s, where s is the shape in question.⁹⁰ But appeal to phenomenal manners of representation will not help us phenomenally distinguish these experiences: V and T both *phenomenally* represent their contents, yet they have different phenomenal properties. There is a difference in what it's like to have them.

If we are persuaded that experiences in different sense-modalities can share content, the right move to make here is to appeal to finer-grained manners of representation: V's relevant avocado shape-related phenomenal property is identical with the impure representational property of visually (phenomenally) representing that something is s, while T's is identical with the impure representational property of tactually (phenomenally) representing that something is s.

As noted above, we can continue this process, appealing to ever finer-grained manners of representation until we have exactly the impure representationist claim we want: perceptually (phenomenally) representing..., visually (perceptually phenomenally) representing..., attentively (visually perceptually phenomenally) representing..., making-the-Necker-Cube-project-outwardly (attentively visually perceptually phenomenally) representing..., and so on.⁹¹

We will do this, though, *only if* we are persuaded that experiences in different sense-modalities can share content, and some might deny this. Chalmers, for example, expresses doubts on this point and allows that the only cases where he sees the need to appeal to finer-grained manners of representation involve Necker Cube-type prominence effect cases (2004a,

⁹⁰ If you do not like this example of a property that can be represented by more than one sense, Block offers the example of the property of being *at that location* (Block forthcoming, p.16).

⁹¹ I do not mean to suggest that every sort of phenomenal difference is best handled by appeal to ever finer-grained specific manners of representation. I just mean to say that the move is available for impure representationists to make.

pp.5-6). Whether we ought to allow that experiences in different modalities can share content is a subtle issue, one I do not want to try to settle here. I will limit myself to the following two observations.⁹²

First, I think it is plausible that there are at least some properties, representation of which is plausibly tied to phenomenal properties, that can be represented by experiences in more than one sense modality. At the very least, this ought to be our default assumption. Surely it is possible.

Second, and most importantly, the choice whether to allow shared content across sense modalities is relevant to the way we taxonomize versions of impure representationism. Let me explain.

Suppose I am an impure representationist who denies that experiences in different sense modalities can share content. Thus, I identify G with the property of phenomenally representing that something is g (= green). Ignoring phenomenal differences due to attention or prominence effects, it follows from this that phenomenal differences *across sense modalities* are determined by representational differences—inter-modal intentionalism. If V and T have no content in common, then, on this view, the phenomenal difference between them is a matter of their phenomenally representing different contents.⁹³

Now suppose that I am an impure representationist who allows experiences in different sense modalities to share content. Thus, I identify G with the property of visually (phenomenally) representing that something is g (= green). Ignoring phenomenal differences

⁹² Note: it isn't a matter of sharing content *per se*. Rather, it is a matter of sharing what we can call, following Chalmers, "phenomenal content," content that we wish to tie closely to phenomenal properties (2004b, p.1). Experiences can have all sorts of content in common, for all the impure representationist cares, as long as that content is not tied closely to phenomenality.

⁹³ Inter-modal intentionalism follows *as long we assume that the two experiences we are talking about both have phenomenality*. If one lacks it, then the phenomenal difference between them will supervene on difference in manner of representation: one will be phenomenal and the other will not be.

due to attention and prominence effects, it follows from this that phenomenal differences *within any given sense modality*—but not across sense modalities—are determined by representational differences—intra-modal intentionalism. If V and T have avocado shape-involving content in common, then, on this view, the phenomenal difference between them is not merely a representational difference.

So whether any particular version of impure representationism is committed to inter- or intra-modal intentionalism will depend on whether the version in question allows that experiences across sense modalities can share content. Thus, we have the possibility of two forms of impure representationism: inter-modal impure representationism and intra-modal impure representationism. For the rest of this chapter, it will not matter for our purposes which form is being discussed.⁹⁴

The impure representationist has many philosophical choices to make. First, there is the choice between inter- and intra-modal forms of the theory. Second, there is the choice between sorts of content: conceptual vs. nonconceptual, propositional vs. Peacockian mixed content (Peacocke 1992), sets of worlds vs. Russellian singular propositions vs. Fregean Thoughts. Third, there is the choice between different account of the required nonrepresentational properties: physical properties, functional properties, or dualist properties. Fourth, there is the choice between different theories of mental representation: causal covariational, teleonymy, etc. Fifth, there is the choice between different accounts of the represented properties: ordinary perceptible properties (e.g. colors, shapes, distances, etc.) or Shoemaker-properties (see below). Sixth, there is the choice between different accounts of the colors: physicalist, primitivist, dispositionalist.

⁹⁴ My guess is that Tye (1995 and 2000) is an inter-modal impure representationist, while Dretske (1995) is an intra-modal impure representationist. But it is very difficult to tell.

One big choice will dictate the answers to many of these smaller choices: the choice whether to be reductionist or not. As we have seen, most impure representationists wish to be reductionists. But this isn't mandatory. There are at least three reductionist/nonreductionist choice points offered by the theory. The impure representationist must give a reductionist account of mental representation (more specifically, of the pure representational and nonrepresentational properties involved in representation) and the represented contents if she wants to give a fully reductionist account of phenomenality.

Dretske (1995a), Lycan (1996), and Tye (1995 and 2000) are all defenses of reductionist impure representationism. Each of them gives a roughly causal covariational account of pure representational properties and a functionalist account of the necessary nonrepresentational properties of experience. It is not clear whether Harman is a pure or impure representationist, but, because of his stress on functionalism, I will interpret him as a reductionist impure representationist (see Harman 1990). Chalmers (2004a) is a rare defense of nonreductionist impure representationism.

Finally, Shoemaker (1994 and 2002) is a reductionist impure representationist, but his case is more complicated. At first blush, Shoemaker looks like a pure representationist because he would identify G with the property—I will not say now whether it is pure or impure—of representing that something is g, where g is, not the avocado's color, but a distinct property Shoemaker calls an "appearance property." (He used to call them "phenomenal properties," but then he changed his usage.)

This looks strikingly like Thau's position. But it isn't. A closer look reveals that Shoemaker's appearance properties are dispositional properties of material objects' surfaces. What is g? The details are complicated—and Shoemaker has changed his mind several times

(see the difference, in particular, between Shoemaker 1994 and 2002)—so let us just say this: *g* is the disposition to cause in certain perceivers an experience with a certain “quale” (Shoemaker’s term). What is important here is that which “qualia” an experience has is, for Shoemaker, determined by functional features of the experience. Thus, Shoemaker’s view turns out to be a slightly more complicated form of reductionist functionalist impure representationism. It is *functionalist* impure representationism because the representational properties the theory appeals to are impure, in virtue of the fact that the properties experiences represent are defined in terms of their capacity to put subjects into certain functional states.⁹⁵

Now we are ready to see whether impure representationism is true.

5. An Inconsistent Triad

The plan for the rest of this chapter is as follows. First, I will show that intentionalism is incompatible with two theses, phenomenal internalism and strong content externalism, both of which are very attractive. Then I will argue that we ought to reject intentionalism so as to maintain the two theses. But, since impure representationism entails intentionalism, impure representationists will have to give up one of the two theses, resulting in either *broad* impure representationism or *narrow* impure representationism. I will argue that we ought to give up broad impure representationism. After that, I will argue that narrow impure representationism is also problematic.

⁹⁵ Two comments. First, a small cottage industry has grown up around Shoemaker’s appearance properties. See Egan (2004) for a survey of the various attempts to puncture and patch Shoemaker’s definitions. Second, why does Shoemaker opt for his theory, rather than one like Tye’s? Because Tye’s is incompatible with the possibility of spectrum inversion without misrepresentation. But Shoemaker can allow spectrum inversion without *color* misrepresentation. He gets to be an intentionalist by saying that color phenomenality supervenes, not on color intentionality, but on appearance property intentionality. By adding a layer of content to experience concerning his appearance properties, Shoemaker is able to be an intentionalist *and* concede something to friends of spectrum inversion cases.

This section shows that intentionalism is incompatible with phenomenal internalism and strong content externalism. Important: it contains extended parts of Egan and John (2004), a paper co-written with Andy Egan.

Three Mutually Incompatible Theses

Intra-modal intentionalism forms an inconsistent triad when combined with the following two theses:

PHENOMENAL INTERNALISM—The phenomenal properties of a mental event supervene on the intrinsic properties of the event's subject.

STRONG CONTENT EXTERNALISM—None of the representational properties of a mental event supervenes on the intrinsic properties of the event's subject.

To see this, consider the following case. Tyler is having a G visual experience, while Andy is having an R visual experience. By phenomenal internalism, Swamp-Tyler is also having a G visual experience, and Swamp-Andy is also having an R visual experience. Since Swamp-Tyler and Swamp-Andy differ phenomenally, by intentionalism, they must differ representationally. But this contradicts strong content externalism, which entails that Swamp-Tyler and Swamp-Andy do *not* differ representationally, since it counts them as being representationally alike in virtue of their both failing to have any representational properties at all).⁹⁶

But if intra-modal intentionalism is incompatible with these two theses, then so is inter-modal intentionalism, since it entails intra-modal intentionalism. But since impure

⁹⁶ Byrne and Hilbert (1997) was the first paper to highlight this inconsistent triad. Chalmers (2004a) presents the inconsistent triad in explicit form.

Note that this argument goes through only if we allow that the version of intra-modal intentionalism in question is one that holds across all possible subjects. A version of intentionalism restricted to experiences within the conscious life of a single subject would avoid the inconsistency. Impure representationists won't like this retreat, though, for it isn't very satisfying as an attempt to support a representational theory of phenomenality. An *intrasubjective* intentionalism would leave possible phenomenal differences between subjects unexplained in

representationism entails intentionalism (of one sort or the other), the impure representationist must abandon one of the two theses to save his theory.⁹⁷ If the impure representationist gives up phenomenal internalism, he will hold *broad* impure representationism: this view identifies phenomenal properties with broad impure representational properties of subjects. (Broad properties fail to supervene exclusively on subjects' intrinsic properties.) If the impure representationist gives up strong content externalism, he will hold *narrow* impure representationism: this view identifies phenomenal properties with narrow impure representational properties. (Narrow properties supervene exclusively on subjects' intrinsic properties.) I will argue that both broad and narrow impure representationism should be rejected.

Why Hold the Theses?

We have already met the reasons—broad theoretical concerns and transparent phenomenology—to hold impure representationism.

Phenomenal internalism is the claim that phenomenal properties are narrow. There are two reasons to believe it.

First, most of us have relatively robust swampman and brain-in-a-vat intuitions that support phenomenal internalism. The feeling is strong that, should a molecule-for-molecule duplicate of me miraculously form in a swamp, the duplicate—Swamp Jim—would be phenomenally the same as me. Likewise for a brain-in-a-vat that is a perfect intrinsic duplicate of my brain.

Second, surely we must look to the brain or some such bodily system to bring about

representational terms. It is consistent with it that for each subject there is a special nonrepresentational quale that all and only experiences had by that subject possess.

⁹⁷ Of course, pure representationism is also inconsistent with the conjunction of phenomenal internalism and strong content externalism.

phenomenal changes. This is the thought that grounds our swampman and brain-in-a-vat intuitions. How could the environment—both spatial *and* temporal—effect a phenomenal change in a subject without effecting a brain or body change? To think that it could would be to think that the spatial and temporal environment can alter the phenomenal properties of experiences by causally bypassing the physical mechanisms making experience possible. It seems extremely plausible that, if you want to change the phenomenal properties of someone’s experience, you have to mess with their brain. Just messing with their environment (in a way that does not induce any changes inside their head) will not do the trick.

This argument moves from the claim (1) that all of the best strategies for changing phenomenal properties involve changing intrinsic properties of subjects’ brains to the claim (2) that phenomenal properties supervene on the intrinsic properties of subjects’ brains. Now, (1) probably does not entail (2), so the argument probably is not deductive. Note, though, that if (1) is true, its truth cannot be a brute fact: there must be *some* explanation of why it is true. The supervenience claim in (2), phenomenal internalism, would nicely explain (1). Moreover, it is difficult to see what non-supervenience-based explanation is available to do the trick. Thus, we have an intuition-based argument to the best explanation for phenomenal internalism.⁹⁸

Strong content externalism is the claim that no content is narrow—that is, that none of the content of a subject’s mental events or states is fixed (entirely) by the intrinsic properties of the subject.

⁹⁸ But isn’t my view that naïve realism is our “pre-theoretical stance” to the world? And isn’t naïve realism phenomenal *externalist*? Why, then, do we have such strong phenomenal internalist intuitions?

Naïve realism, which is phenomenal externalist, is our pre-theoretical stance insofar as, were we to spell out in words (as I have tried to do) the way phenomenally conscious experience strikes us, we would find that naïve realism results. This does *not* mean either that naïve realism is our folk theory of experience or that we can’t also have strong phenomenal internalist intuitions. (Us scientifically informed moderns tend to have extremely powerful phenomenal internalist intuitions. Professions of faith in phenomenal externalism, as phenomenal externalists such as Dretske and Lycan are well aware, tend to be greeted with incredulous stares.) I think that part of what makes the

There is an extensive literature motivating externalism about content.⁹⁹ It is worth noting, though, the difference between the weaker claim that some content is broad from the stronger claim discussed here, that all content is broad. Arguments that focus on particular kinds of content, like those involving natural kinds, will not be enough to establish the strong claim.

Swampman cases might do the trick, though. It is quite plausible (at least once we have bought into a broadly causal theory of content) that Swampman's experiences do not have any content at all, since Swampman is not connected to his environment in the right way. If this is right, then no content can be narrow, since Swampman's experiences do not have any content in common with his intrinsic duplicate with a normal history and origin.

Another motivation (perhaps the one that is behind the swampman intuitions) is that representation is to be explained in broadly causal terms—the representational properties of states depend on their causal connections to various things in, and features of, the subject's environment. Take away all the causal connections, and you take away all the content. There is no aspect of the content that cannot be changed or removed by changing the way that the subject is embedded in their environment.

6. Give Up Broad Impure Representationism

So intentionalism is incompatible with phenomenal internalism and strong content externalism. Note, though, that phenomenal internalism and strong content externalism do not enjoy the same degree of intuitive support. While I find denials of strong content externalism difficult to understand (how could a subject's intrinsic properties make it the case that one of her mental

counterintuitive discoveries of vision science so thrilling—sometimes even frightening—is the clash between the naïve realism-inspiring phenomenology of experience and our science-supported phenomenal internalist intuitions.⁹⁹ The classic works are Putnam (1975) and Burge (1979). It is worth noting that Putnam and Burge seem to have assumed phenomenal internalism.

events or states has a truth condition?), denials of phenomenal internalism seem beyond the philosophical pale. And even those philosophers who are less certain of strong content externalism feel more confident, I'm sure, in phenomenal internalism than in any thesis about content they might hold.¹⁰⁰

This means, I think, that we should look more favorably on the narrow impure representationist response to the inconsistent triad than on the broad impure representationist response. Phenomenal internalism really does seem non-negotiable; thus, broad impure representationism has to go.

But if broad impure representationism has to go, then most extant forms of impure representationism have to go, for Dretske, Lycan, and Tye are all staunch content externalists and proud—if beleaguered—defenders of phenomenal externalism. (I am less sure of Harman.) For the rest of this dissertation, then, I will assume that we can rule out broad impure representationism.

I'll bet you're thinking that this is the most outrageously dogmatic dismissal of a philosophical position you have ever encountered. Let me try to address this very serious worry.

To begin with, as I have said, I really don't think broad impure representationism is very plausible—its commitment to phenomenal externalism seems fairly close to a *reductio*.

Having said this, though, I must point out that broad impure representationists have recently begun addressing this charge. Dretske and Lycan, for example, have both written papers defending their phenomenal externalism (see Dretske 1995b and Lycan 2001). I plan to address their arguments in future work.¹⁰¹ (I want to stress that I mean no disrespect to Dretske,

¹⁰⁰ I find that I can get my mind around the very idea of phenomenal externalism only by assimilating it to naïve realism.

¹⁰¹ I also hope to address in future work another oversight. One might think, in light of Byrne's argument *for* intentionalism, that all I've shown is that there is a paradox of phenomenality (see Byrne 2001). On one hand, we

Lycan, and Tye when I set aside their views—I don't do so lightly. Tye, in particular, has an impressively detailed and very powerful representationalist theory that, in my opinion, is liable to be the most worked-out philosophical account of phenomenality we see for awhile.)

I am especially sensitive to the charge of dogmatism since the theory I will defend, a form of pure phenomenism, has a consequence that many will regard as *at least* as implausible as broad impure representationism's phenomenal externalism. I think there is a big problem here. Indeed, the existence of phenomenality poses *two* really big problems: one is the problem of the explanatory gap and the other is the inconsistent triad. I think the most plausible theories of phenomenality turn out to be either some form of broad impure representationism or Shoemaker's narrow impure representationism (see below) or my pure phenomenism (see the next chapter). But all three have tremendous problems. If my dissertation does one thing it will be charting in a bit more detail the contours of this second puzzle of phenomenality.

At all events, we have decided to rule out broad impure representationism. So let us turn now to the question of how plausible narrow impure representationism is.

7. Give Up Narrow Impure Representationism

There are many ways in which narrow impure representationism might be developed. I will consider only two of them. The first, Shoemaker's theory, is Russellian. The second, David Chalmers's theory, is Fregean.

Shoemaker's Theory

have Byrne's apparently sound argument for intentionalism; but, on the other hand, we have intentionalism's incompatibility with two seemingly unimpeachable theses. (Byrne, for his part, is more certain of intentionalism than he is of either phenomenal internalism or strong content externalism.)

We have already met Shoemaker's theory. It is a form of narrow impure representationism because Shoemaker believes that appearance property content is determined by subjects' intrinsic properties. His theory is Russellian: the appearance property contents to which Shoemaker has us related in perceptual experience are usefully regarded as Russellian singular propositions, containing appearance properties as constituents.

Shoemaker's theory is very attractive. By appealing to two tiers of perceptual content—one, color (and shape, distance, etc.) content, determined externally, and the other, appearance property content associated with color (and shape, distance, etc.) content, determined internally—Shoemaker combines intentionalism (about appearance property content) and the possibility of spectrum inversion without (color) misrepresentation. Moreover, his view is attractively reductionist in that it gives us a plausible functionalist account of the nonrepresentational "qualia" it demands.

But it has two problems. First, specifying exactly what the appearance properties are is very difficult. See Egan (2004) for discussion of the difficulties. Second, the appeal to appearance properties themselves is problematic. Even if we could solve the technical problems besetting Shoemaker's commitment to appearance properties, their very existence poses problems. Surely the salient surface property *g* of the avocado you see in having your visual experience *E* is the avocado's *color*.

This complaint actually runs together two distinct worries. One is error theoretic: Shoemaker's theory entails that a lot of our perceptual-phenomenal judgements are wrong. When you point to the avocado's surface and say "*that* property, the one chiefly responsible for whatever is greenish or green-feeling about my experience, is green" you are wrong: the avocado's color, on Shoemaker's theory, isn't that property. The other worry is ontological:

Shoemaker's theory dulls Occam's Razor. If possible, we should try to avoid positing new properties.

Neither of these criticisms comes even close to being a knockdown objection to Shoemaker's theory. What they do accomplish, though, is to show that Shoemaker's theory ought to be a last-resort position. So let us set it aside.

(I do not set this view aside lightly. I wish I had more to say against it. But consider the following. Shoemaker is a very strong proponent of phenomenal internalism and a strongish proponent of intentionalism. This is why his solution to the inconsistent triad is to give up strong content externalism. Note, though, that he seems to hold intentionalism only because of a belief that the transparency of experience uniquely supports it. But as I will try to show in the next chapter, this is wrong: neither pure nor impure phenomenism is in conflict with transparency.

Chalmers's Theory

Chalmers (2004a) defends a form of nonreductionist narrow impure representationism. It is nonreductionist because he thinks that irreducibly phenomenal notions must be employed in specifying the representational properties the theory calls for. It is also a Fregean theory: the representational properties it requires are Fregean. Let me explain.

Chalmers writes: "Let us say that a *Fregean representational property* is the property of having a certain Fregean content (in a certain way). Let us say that *Fregean representationalism* is the thesis that phenomenal properties are [identical with] certain (pure or impure) Fregean representational properties" (p.14). Take your visual experience E* of a pomegranate. It is, in our terminology, an R visual experience. Chalmers (rather confusingly) introduces the term 'phenomenal redness' to stand for the same property our 'R' stands for. (Chalmers is clear that

phenomenal redness is the same property as R; in particular, phenomenal redness is a property of E*, not of the pomegranate you see.)

Then he writes: “one might propose that phenomenal redness is [identical with] the property of having a certain Fregean content (in the appropriate phenomenal way), where this Fregean content involves a mode of presentation such as *the property that normally causes experiences of phenomenal redness*” (ibid, emphasis Chalmers’s). He goes on:

On this view, the relevant representational content does not directly involve the property attributed by the experience. It may well be that the experience attributes the property of redness [not phenomenal redness] to an object, and that redness is a surface spectral reflectance property. This attributed property may enter into the Russellian content of the experience, but it does not enter into the Fregean content. Rather, the Fregean content involves a mode of presentation of this property (ibid).

Thus, Chalmers’s view is that R is the property of being a mental event with the impure representational property of representing the Fregean content—call it “C”—in a phenomenal manner, where C “involves” a mode of presentation of redness. This is an attractive position. Chalmers can be an intentionalist of sorts (phenomenal properties supervene on Fregean, not Russellian, representational properties), accept phenomenal internalism (Fregean representational properties, he thinks, are internally determined), accept strong (Russellian) content externalism (he thinks that experiences’ Russellian contents are externally determined), and allow spectrum inversion without (Russellian) misrepresentation (since two experiences can have the same Russellian content under different Fregean modes of presentation).¹⁰²

There is a lot to say about this position. I will confine my comments to one issue: what

¹⁰² As Chalmers notes, his theory is closely related to Shoemaker’s. He writes: “Both views give a key role to dispositional notions such as *normally causes phenomenally red experiences*. The difference is that where Shoemaker’s view holds that color experiences attribute these dispositions to objects, my view holds that color experiences attribute intrinsic properties (colors) to objects, with the dispositional notions serving as modes of presentation of these intrinsic properties” (p.15). This difference, he thinks, corrects a flaw in Shoemaker’s theory: “[On my view] a dispositional relation to experience is used to *determine* the property attributed by the experience, but the property attributed is not itself dispositional. This avoids one major difficulty for Shoemaker’s view: the claim that the primary properties attributed by color experiences are not colors” (ibid).

is C? Chalmers says that it is a Fregean representational content. To explain what he means by this I have to say a bit about Chalmers's treatment of linguistic and belief content.

Chalmers, like Frege before him, says that linguistic expressions have a sense and a reference. The reference of 'Hesperus' is the planet Venus; its sense is a mode of presentation of that planet. "Property terms" (p.12) also have sense and reference. The reference of 'bright' is the property of brightness; its sense is a mode of presentation of brightness.

Chalmers extends this account to mental states:

A belief such as *Hesperus is bright* is composed of concepts, such as *Hesperus* and *bright*. The concepts can be said to have extensions: the planet Venus and the property of brightness, respectively. The concepts can also be said to have *modes of presentation* of those extensions. These modes of presentation have the same extension but different modes of presentation. The belief as a whole may have a complex content that is composed of these modes of presentation (ibid).

His answer to the question "What is a mode of presentation" goes as follows: "...one natural approach characterizes a mode of presentation as a *condition on extension*. The idea is that every concept is associated with some *condition*, such that an entity in the world must satisfy this condition to qualify as the extension of the concept" (ibid, emphasis Chalmers's).¹⁰³

Note something important. In the linguistic case, we are able intelligibly to assign senses and references to parts of sentences only because, prior to the assigning, we are provided with uninterpreted sentences of a language with a well-defined syntax. The syntax carves up the sentences into syntactic atoms—singular terms and predicates—to which we can then assign distinct senses and references. Without the well-defined syntax, we would not know to which bits of the sentence to assign which senses and references. Fregean semantics requires an antecedent syntax.

Now consider the mental case. Chalmers is able to extend his Fregean treatment to belief

¹⁰³ Chalmers says much more than this in other places. His "conditions on extension" turn out to have a complicated modal structure (see Chalmers 2002).

and desire only by assuming that beliefs are “composed of concepts.” These “concepts” play the role played by singular terms and predicates in the linguistic case. But beliefs are *properties of subjects*; as such, they have no parts—they aren’t composed of anything. To make sense of this assumption, we have to interpret Chalmers as thinking that when a subject has a belief she tokens something with sentence-like structure and that there is a well-defined syntax for such tokens. This, I think, is the most plausible reading of Chalmers’s claim that beliefs are “composed of concepts” that have extensions and intensions. It attributes something close to, if not the same thing as, Fodor’s Language of Thought theory (see Fodor 1975) to Chalmers. This is a significant further assumption. To make the attribution of senses and references to beliefs work the way he wants it to work, Chalmers assumes that beliefs either are or involve sentence-like tokens with an antecedently determined syntax. With this assumption in place, the assignment of senses and references to the various “parts” of beliefs can proceed.¹⁰⁴

But now consider what he says about the case of perceptual content:

One can extend this approach to the content of perceptual experiences. Perceptual experiences attribute properties to objects: e.g., my visual experience might attribute greenness to a ball. And as with beliefs and concepts, we can hold that a perceptual experience involves some *modes of presentation* of these properties and objects. As before, these modes of presentation can be seen as conditions on extension. For example, there is a condition that an object must satisfy in order to be the object represented by my experience. The ball satisfies this condition, so it is the object represented by my experience. There is also a condition that a property must satisfy in order to be the property attributed by my experience. Greenness satisfies this condition, so greenness is the property represented by my experience (p.13, emphasis Chalmers’s).

This talk of experiential conditions on extension sounds so plausible that it almost slips by without notice until one realizes that something is conspicuously absent from this passage. When Chalmers talks about perceptual content he mentions nothing that could serve as perceptual analogs of the sentence in the linguistic case. There were mental analogs of the

¹⁰⁴ I owe these observations to Robert Stalnaker. He expressed them in his Fall 2001 MIT graduate seminar on mental representation.

sentence in the belief case: the items, presumably mental tokens instantiated in belief, composed of “concepts.” But Chalmers does not say that perceptual experiences are composed of “concepts” as he says beliefs are. He doesn’t say they are composed of anything.

Here is the worry posed by the absence of sentence analogs. Remember: Fregean semantics works only if there is something sentence-like with a well-defined syntax it can *work on*. In the linguistic case, the sentence-like things were just sentences. We know how to assign senses (modes of presentation) and references to parts of sentences because we have a thing—a sentence—and a theory—syntax—telling us how to divide the thing up into bits (singular terms and predicates). These bits can then be assigned senses and references. Such a story works for the belief case only if there is a sentence-like thing—presumably a sentence of mentalese—and a theory—presumably the syntactic theory yielded by completed Language of Thought science—telling us how to divide the thing up into bits (“concepts”). Only then can we speak of beliefs involving senses and references. So if a similar story is to work in the perceptual case, we need, corresponding to your experience E* of the pomegranate, a thing—a sentence of “perceptionese” —and a theory—the syntactic theory yielded by completed Language of Perception science?—telling us how to divide the thing up into bits. Only then can we speak of perceptual experiences involving senses and references.

I think this poses a dilemma for Chalmers. Without something like a Language of Perception, he has no way of coherently and consistently assigning modes of presentation to experiences, but the very idea of a Language of Perception is far-fetched. Fodor and others appeal to a Language of Thought chiefly to explain the way in which thought seems, as language does, productive and systematic. It is not clear that there is any sense in which perceptual experience exhibits either productivity or systematicity. Chalmers has implausibly *linguisticized*

perception and phenomenality.

I submit that Chalmers's Fregean narrow impure representationism is workable only if we assume that something resembling the Language of Thought Hypothesis is true for phenomenally conscious perceptual experience. I strongly doubt that this is the case. Therefore, we ought to reject Chalmers's theory.

8. Conclusion

I have explained impure representationism, shown that it is incompatible with the conjunction of phenomenal internalism and strong content externalism, and argued that broad and narrow impure representationism are untenable.

Chapter 10: Pure and Impure Phenomenism

1. Introduction

We have seen that pure and impure representationism are untenable. We had better abandon the hope of explaining phenomenality in intentional terms and opt for either pure or impure phenomenism.

In this chapter, I defend pure phenomenism. After showing that impure phenomenism suffers from difficulties similar to those suffered by impure phenomenism, I will explain pure phenomenism. Then I will defend it against the representationist charge that it conflicts with the transparency of experience. But my defense of pure phenomenism on this point will bring to light features of my particular version of the position that many might regard as highly problematic. I will argue that these features aren't really problematic at all.

The chapter concludes with a brief summary of this dissertation's main results and a brief prospectus for further research.

2. Impure Phenomenism

One reaction to impure representationism's troubles is to think that what is wrong with the theory is its commitment to intentionalism, either inter- or intra-modal. The philosopher who reacts this way might fall back on impure phenomenism.

What is Impure Phenomenism?

Recall that impure phenomenism says that phenomenal properties like G are identical with complex properties composed of nonrepresentational properties of experience together with

some pure representational property or other. Impure phenomenism often refers to its favored nonrepresentational properties as “qualia.” Thus, it entails that G is identical with the property of having a certain quale and having some content or other.

This theory is *impure* phenomenism because it makes a mental event’s having some content or other necessary for it to have phenomenality. It is *impure phenomenism* because it denies intentionalism, inter- or intra-modal. So impure phenomenism constitutes a kind of retreat from impure representationism. It still makes content necessary for phenomenality; but it lets pure representational properties and phenomenal properties come apart. Thus, it is compatible with impure phenomenism, but not with impure representationism, that a visual experience is G yet it represents the content that something is r or that a visual experience is R yet it represents the content that something is g. So impure phenomenism is compatible with the possibility of spectrum inversion without misrepresentation.¹⁰⁵

Impure phenomenism, as I have just presented it, leaves open the answer to the reduction question. For all I have said, its qualia might be irreducibly phenomenal. Or they might be physical or functional properties of subjects. Likewise, I have left open whether it gives a reductive or nonreductive account of intentionality and the properties represented by experience. Though all of these questions are open, nothing in what follows will turn on how they are best closed.

Impure phenomenists can allow that perceptual experiences are propositional attitudes, but, since they reject intentionalism, they cannot allow that experiences are individuated by their contents alone. They will choose to individuate experiences by reference to their contents *and*

¹⁰⁵ Though impure phenomenism says that G is identical with the property of having a certain quale and having some content or other, it doesn’t follow that *any* content or other is sufficient, when conjoined with the quale in question, for G. For example, impure phenomenists do not have to allow that an experience’s having *shape* content and our quale entails its having G.

their qualia, with different impure phenomenists offering different accounts of how to divide the labor of individuation between the two. At all events, they have no trouble giving a direct realist answer to the problem of direct awareness. That is because, as an application of the Possession Model, impure phenomenism says that all subjects are aware of in having their experiences are the objects and properties their experiences are experiences of.

Against Impure Phenomenism

The biggest problem with impure phenomenism is that it doesn't help: retreating to it from impure representationism does not resolve the inconsistent triad. For there is a *new* inconsistent triad formed by impure phenomenism, phenomenal internalism, and strong content externalism.

Impure phenomenism entails the following:

CONTENTISM—A mental event can have a phenomenal property only if it has at least some content.

Contentism is the weakest possible way of tying phenomenality to intentionality. (Strictly speaking, pure and impure representationism also entail it. It is just that they also entail intentionalism, whereas impure phenomenism does not.) If you deny contentism while maintaining allegiance to the Possession Model, you have got to take up pure phenomenism.

Now consider Jones. He is having an R visual experience. By phenomenal internalism, Swamp Jones is also having an R visual experience. By strong content externalism, Swamp Jones's experience has no representational properties—no content. Thus, Swamp Jones is enjoying a contentless phenomenally conscious mental event. This contradicts contentism.¹⁰⁶ Something must go. Since phenomenal internalism and strong content externalism, I contend,

¹⁰⁶ Essentially the same argument appears in Block (forthcoming), p.6.

are mandatory, we should give up contentism.

Note carefully: in saying that we should let go of contentism, I am *not* saying that phenomenally conscious perceptual experiences lack locationality. In denying contentism, I do not thereby deny the very phenomenology of experience. I merely deny that an experience's having content—being such as to determine a truth-condition—is necessary for it to have locationality.

Are there any impure phenomenists? It is difficult to say. The most vocal opponent of representationism in the philosophical literature is Ned Block. He has presented a battery of arguments—an array of inverted spectrum arguments, including, most famously, his “inverted earth” argument (see his 1990)—designed to show that representationism is false. But it is sometimes difficult to tell whether he holds impure phenomenism or pure phenomenism. (He calls himself a “phenomenist.”) In his paper “Inverted Earth” (1990), he speaks of “phenomenal content,” giving the impression that content is necessary for phenomenality. But in other places, particularly in his recent “Mental Paint” (forthcoming), he writes as if content is not necessary for phenomenality, denying that orgasms, for example, have representational content and presenting a version of the above Swamp Jones argument. At all events, if Block holds impure phenomenism, he should not, for it is incompatible with phenomenal internalism and strong content externalism, claims to which he is very strongly attached (indeed, they drive his inverted earth argument).

3. Pure Phenomenism

So we are left with pure phenomenism. It is the only application of the Possession Model compatible with both phenomenal internalism and strong content externalism. I conclude that

phenomenality cannot be explained in terms of intentionality.¹⁰⁷

Pure phenomenism says that phenomenal properties like G are identical with nonrepresentational properties of experience. It calls these nonrepresentational properties “qualia.” (This is why I said in the first chapter that one can deny the existence of qualia without denying the existence of phenomenal properties.) Thus, pure phenomenism entails that G is identical with a certain quale.

Pure phenomenism is compatible with thinking that experiences have representational content. It simply denies that content is necessary for phenomenality.¹⁰⁸ I am not sure what pure phenomenists should think experiences are. They probably ought to say that “experience” talk is shorthand for talk of subjects’ instantiating qualia. If a subject instantiates a quale and also has some pure representational properties, then the subject has an “experience” with phenomenality that represents the world as being a certain way.¹⁰⁹

Don’t think that pure phenomenism must give a dualist nonreductionist answer to the reduction question. For all I have said, the answer to that question is wide open. Pure

¹⁰⁷ Once more, I want to emphasize that I am aware of the fact that my dismissals of impure representationism and impure phenomenism have been quick. But, even if you are utterly unconvinced, you can still regard this final chapter as mapping out a bit of logical space that hasn’t been mapped out before. After all, what *does* follow when we put together, as I am about to do, phenomenal internalism, strong content externalism, pure phenomenism, and the claim that the only properties subjects can be aware of in having their experiences are the properties (if any) their experiences represent?

¹⁰⁸ Throughout this dissertation I’ve described qualia as nonrepresentational properties. What I ought to say is that they aren’t *intrinsically* or *inherently* representational. This qualification is needed because pure phenomenism can allow that experiences can—and almost always do—have representational content. But if the pure phenomenist says that experiences are just instantiations of qualia by subjects, then qualia *are* representational, at least insofar as they have representational content. What’s really important to pure phenomenism is that (1) intentionalism and its converse are false and (2) contentism is false. Block captures this in slogan form by saying that phenomenality “outruns” intentionality (Block, forthcoming, p.1).

¹⁰⁹ Two points.

First, nothing forces a pure phenomenist to accept the view of experiences that I suggest. They could reserve the term ‘experience’ to apply only to cases in which subjects instantiate qualia and (perceptual) representational properties. On this view, even though Jones and Swamp Jones are phenomenally identical, only Jones enjoys an “experience” since only he is undergoing a mental event with content.

Second, I introduced the term ‘qualia’ as a name for phenomenality-determining nonrepresentational properties of *experiences*, but here I am using the term as a name for phenomenality-determining

phenomenism is compatible with nonreductionism (on which qualia are irreducibly phenomenal properties) and with reductionism (on which qualia are either physical or functional properties). The version of pure phenomenism I prefer says that qualia are physical properties of subjects' brains and perceptual systems.¹¹⁰

As I noted above, Block seems to defend pure phenomenism in at least some places. See Block (forthcoming). Levine (2001) also seems to defend it. But neither philosopher is particularly clear on the matter.¹¹¹

My theory of phenomenality, then, combines pure phenomenism, phenomenal internalism, and strong content externalism. I think this provides us with a plausible account of phenomenality. On my view, perception works as follows. As a result of causal perceptual interactions with instantiations of perceptible properties (colors, shapes, distances, etc.), subjects undergo mental events that instantiate certain properties. Among these properties are intrinsic, nonrepresentational properties, call them "qualia," which causally covary with distal instantiations of perceptible properties. My theory identifies phenomenal properties with these qualia. Qualia are inner proximal "signs," if you like, of outer distal perceptible properties. (Be careful, though: pure phenomenism's qualia are not properties subjects are aware of in having their experiences, as sense-datum properties are according to the sense-datum theory—see below.) Qualia—the phenomenal properties—determine what it's like for subjects to have their experiences.

What goes on in the case of spectrum inversion without misrepresentation, according to

nonrepresentational properties of *subjects*. Nothing of real interest turns on whether we take qualia to be properties of subjects or of subjects' experiences. These are just different ways of keeping the books.

¹¹⁰ I realize that this poses problems for my view in light of the alleged multiple realizability of phenomenal properties. I won't have anything to say about this.

¹¹¹ Chalmers (1996) seems to presuppose pure phenomenism. But whether it does or not, Chalmers has since changed his mind.

my view? Simple. Suppose Norm is not spectrum inverted, while Abnorm is. Norm and Abnorm are looking at a green avocado. Neither misrepresents its color: both represent the avocado as green.¹¹² But Norm's inner "sign" for outer green is the quale I identify with the phenomenal property G, while Abnorm's inner "sign" for outer green is the quale I identify with the phenomenal property R. If you like, Norm and Abnorm have different ways of perceptually "getting at" or, perhaps better, "tracking" the color green: Norm in the G way, Abnorm in the R way. (Again, these qualia—G and R—are *not* properties of which Norm and Abnorm are aware in having their experiences—see below.)

4. Pure Phenomenism and Transparency

Representationists often argue that theories which appeal to qualia conflict with transparency.

This section contends that this argument from transparency is unsound.

Transparency

In chapter two, I introduced the notion of transparency as follows.

First, I pointed to the presentationality and locationality of phenomenally conscious perceptual experience.

If you introspect your experience E of the avocado, you will find that there is something about what it's like to undergo it in virtue of which you are disposed to judge that E makes g present to your mind. There is just something about what it's like to have E in virtue of which you are spontaneously inclined to report things like: "When I have E, g is somehow *present to*

¹¹² It is plausible that Abnorm does not misrepresent the avocado's color. His R experiences allow him to track green objects as well as Norm's G experiences allow him to track them. Abnorm is a competent English speaker; he means by 'green' what Norm means by it. And it would violate the principle of charity to suppose that Abnorm's experiences are systematically nonveridical. See, however, Hilbert and Kalderon (2000) for an opposing view.

my mind, right there, before me—all in a way it isn't when, with eyes closed, I think thoughts about *g* things." This "something" is E's presentationality.

If you introspect E, you will also find that there is something about what it's like to undergo it in virtue of which you are disposed to judge that, for every property you take it to make present to your mind, E locates the property in your external, mind-independent environment. There is just something about what it's like to have E in virtue of which you are spontaneously inclined to report things like: "When I have E a property is present to my mind but also *located* by my experience in my environment. It's as if in having it E *says* that *g* is instantiated. It's as if there is a way the world is *according to* my experience." This "something" is E's locationality.

Second, I claimed that an experience is transparent if and only if (1) it has presentationality and locationality and (2) introspection of it reveals that, in having the experience, the only properties the subject is aware of are the properties she is disposed to judge are presented and located.

Finally, I pointed out that there are four ways for a theory of phenomenality to conflict with transparency, only one of which is significant for our purposes: a theory of phenomenality could entail that the properties subjects judge to be presented and located by their experiences are not instantiated where their experiences locate them. We should respect the transparency constraint and reject theories of phenomenality that conflict with transparency.

The Argument from Transparency

To keep things simple, let us say that a property is a "presented property" if it is a property of which a subject is spontaneously inclined to judge that it is present to mind in

experience. *g*, then, is a presented property. The argument from transparency runs as follows:

THE ARGUMENT FROM TRANSPARENCY

- P1 If pure phenomenism is true, then the presented properties are properties of experiences.
- P2 The presented properties are not properties of experiences.
- C Pure phenomenism is false.

According to P1, pure phenomenism has the consequence that the properties subjects are disposed to judge are presented are instantiated by experiences. According to P2, that's false: the presented properties do not belong to experiences. To think otherwise would be to conflict with the transparency of experience, for experiences are such that in virtue of what it's like to have them subjects are disposed to judge that the presented properties are located by their experiences in the external, mind-independent environment. Therefore, pure phenomenism is false.¹¹³

The argument from transparency is unsound.

There are at least two different ways to put qualia—remember: nonrepresentational properties of experience—to work in an explanation of phenomenal properties. To see this, set aside for the moment what I have said about pure phenomenism and consider the following two “qualia theories” of phenomenal properties:

QUALIA POSSESSION (QP)—An experience has a phenomenal property only if it has a quale. This is because phenomenal properties are qualia.

QUALIA CONFRONTATION (QC)—An experience has a phenomenal property only if (1) it has a quale and (2) the experience's subject is brutally confronted with the quale in having the experience. This is because phenomenal

¹¹³ Representationalists tend to present the Argument from Transparency in an impressionistic way. I believe that my reconstruction captures their major aim (or, at least, one of their major aims) in running the argument. In fact, what I offer is basically a restatement and simplification of part of Tye (2002). Martin (2002) discusses the argument in something like the form I give it. Stoljar (forthcoming) presents a different reconstruction of the argument.

properties are complex properties composed of qualia and relations of brute confrontation with qualia.

These claims require some explaining.

QC is an application of the Confrontation Model. What it is for E to have G, it says, is for you to be brutally confronted with a quale in having E. In other words, E's having G is constituted by your brutally confronting a quale.

QP is an application of the Possession Model. It says that your experience has G because it possesses a quale, not because you are brutally confronted with a quale in having E. QP identifies phenomenal properties with qualia.

Advocates of the Confrontation Model believe that the properties subjects are brutally confronted with in having their experiences are the presented properties. Thus, Sense-Datum Theorists argue that g is identical with the sense-datum property their theory says you are brutally confronted with in having E. Likewise, philosophers who endorse QC argue that g is identical with the quale their theory says you are brutally confronted with in having E.

Advocates of the Possession Model believe that the presented properties are the perceptible properties (colors, shapes, distances, etc.) subjects' experiences represent things in their worldly environments as instantiating, properties which belong not to sense-data or experiences but to material objects (or nothing, as is the case with hallucinations). They do not believe that the special properties their theories call for—depending on the theory, pure representational properties, impure representational properties, or qualia—are presented properties. Thus, pure representationists, for example, argue that g is a property of the avocado you see. (Most think that g is green.) Philosophers who endorse QP also argue that g is a property of the avocado you see. (Most also think that g is green.)

Now we can see why the distinction between QC and QP matters to the argument from

transparency. Recall that P1 says: if pure phenomenism is true, then the presented properties are properties of experiences. And P2 says: the presented properties are not properties of experiences. P1 is true only of QC, not of QP. To escape the argument, friends of QC must deny P2, and that is to conflict with the transparency of experience. But friends of QP can happily grant P2, rejecting P1 instead. So pure phenomenists have nothing to fear from the argument from transparency provided they accept QP.

But pure phenomenists do hold QP: pure phenomenism *is* QP. Representationalist critics of pure phenomenism, though, almost always assume that pure phenomenism is QC. But I am unable to think of a single contemporary philosopher who defends QC. For that matter, I can't think of anyone who has ever defended QC.¹¹⁴ This is for the best, since QC is an application of the Confrontation Model with all of the attendant problems.

Some might protest that pure phenomenism/QP is incoherent. "How," they might ask, "could properties we are not aware of in having our experiences—qualia, as pure phenomenism/QP conceives of them—explain phenomenal properties?" Representationalists, of course, have no business making this objection. For if it works against pure phenomenism/QP, it works against pure and impure phenomenism. They also explain phenomenal properties by identifying them with properties we are not aware of in having our experiences.

In any case, one will find this objection compelling only if one is wedded to the Confrontation Model. After all, *what it is* for an experience to have a phenomenal property, according to the Confrontation Model, is for the experience's subject to be brutally confronted with a property in having the experience, and brute confrontation is, at the very least, a species of

¹¹⁴ In fairness to representationalists, I ought to note that pure phenomenists themselves are not always careful to highlight the QP/QC distinction. I suspect that this is because pure phenomenism is usually discussed in the consciousness literature, where careful attention to the structure of experience is not as important as it is in the perception literature.

awareness.

But the Confrontation Model is not mandatory. As we have seen, the Possession Model is preferable. And, in any case, there simply is not anything odd—let alone incoherent—about pure phenomenism/QP’s way of explaining phenomenal properties. We are all familiar with the way in which the light in a room—the light itself, not the light source—affects what it’s like for us to have experiences of things in the room. If someone were to tamper with the market’s lighting, altering the light inside, what it’s like for you to see the avocado would change. The light in the market can condition what it’s like for you to see the avocado even if you are not aware of it in seeing the avocado. For pure phenomenism/QP, qualia are a bit like the light in a room. They condition what it is like for us to have our experiences without our having to be aware of them in having those experiences.¹¹⁵ There is nothing incoherent here.

The Argument from Transparency goes through, then, only if it is assumed that pure phenomenism is QC. But pure phenomenism isn’t QC. It is QP.

So far, I have defended the combination of pure phenomenism, phenomenal internalism, and strong content externalism. But, in light of this discussion of transparency, I want to add a further claim to my package of views, the “Awareness Principle”:

AWARENESS PRINCIPLE—All subjects are aware of in having their experiences are the objects (if any) they perceive and the properties (if any) their experiences represent those objects as having.

All you are visually aware of in having E, according to the Awareness Principle, is the avocado you see and the properties you visually represent it as having. If E had been a hallucination, you would have been visually aware of no object in having it. If you had been like Swamp Jones,

¹¹⁵ Be careful: qualia, according to pure phenomenism/QP, are only *a bit* like the light in a room. One important difference is that, while it might be possible to become aware in having E of the light in the market, it isn’t possible to become aware in having E of E’s qualia. (That isn’t to say, though, that you can’t be introspectively aware of E’s qualia—see below.)

you would have been visually aware of no properties in having it.

Here is why I hold the Awareness Principle. I think that the properties subjects are most basically aware of in having their experiences are the properties (if any) they judge to be presented and located by their experiences. I also think that the properties subjects are most basically aware of in having their experiences are the properties (if any) subjects' experiences represent their environments as instantiating. By holding the Awareness Principle, I can say that the only properties subjects are aware of in having their experiences are properties that belong, if they belong to anything at all, to material objects in subjects' environments. The property *g* you are inclined to describe as present to your mind in having *E* is a property that your experience represents your environment as instantiating, and *g* is among the properties you are most basically aware of in having *E*. In having *E*, there is nothing but the avocado and the properties *E* represents it as having for you to be aware of. Commitment to the Awareness Principle flows from a desire to respect the transparency constraint.

5. The No Awareness Problem

In this section, I show that my view of phenomenality leads to two consequences that some might consider highly counterintuitive. I argue that, while they may be counterintuitive, the consequences aren't grounds for rejecting my view.

Remember: my view is pure phenomenism, phenomenal internalism, strong content externalism, and the Awareness Principle.

Black-and-White Maria

Maria was raised in a black and white environment. She was never told about the colors.

Indeed, she has had no causal contact at all with colors. Not only that, but she has never even had an experience—whether caused by a colored object or not—with the phenomenality we associate with color experience.¹¹⁶

Maria, then, is with respect to color content the way Swampman is with respect to all content: she doesn't have any.

One day Maria is in the black and white laboratory of a neurosurgeon. He puts Maria into a deep, dreamless sleep. Then, at time t_1 , he wakes her up while accidentally stimulating her brain to go into brain state 1. After that, at time t_2 , he accidentally stimulates her brain to go into brain state 2.

Since we are assuming pure phenomenism and phenomenal internalism, we can assume that brain state 1 is sufficient to bring about a B visual experience in Maria—as it happens, the sort of experience she would undergo if she were to see something blue. We can also assume that brain state 2 is sufficient to bring about a Y visual experience in Maria—as it happens, the sort of experience she would undergo if she were to see something yellow.

At t_1 Maria concentrates intently in having her B experience. If you were in Maria's place, it would seem to you as if your visual field had gone totally blue, as if you were standing before a giant, uniformly illuminated blue wall. At t_2 she concentrates intently in having her Y experience. If you were in Maria's place now, it would seem to you as if your visual field had gone totally yellow, as if you were standing before a giant, uniformly illuminated yellow wall. Let us stipulate that Maria notices the phenomenal change from having a B experience at t_1 to having a Y experience at t_2 . Never having seen any colors, she's amazed by it all.

Now, my pure phenomenism says that phenomenal properties cannot be explained in terms of representational properties, and it allows that Swamp Jones or a subject like Maria could

¹¹⁶ Assume that Maria's acts of imagining and dreams also lack such phenomenality.

have “color” phenomenality without any color content (because of its dual commitment to phenomenal internalism and strong content externalism). Moreover, it holds the Awareness Principle.

The First Counterintuitive Consequence

At t_1 , by content externalism, Maria’s experience lacks blue content. There is no object she is visually aware of in having her B experience. Since her experience lacks blue content, it does not represent blue; thus, by the Awareness Principle, there is no property she is visually aware of in having her experience. At t_2 , by content externalism, Maria’s experience lacks yellow content. There is no object she is visually aware of in having her Y experience. Since her experience lacks yellow content, it does not represent yellow; thus, by the Awareness Principle, there is no property she is visually aware of in having her experience.

Given the Awareness Principle, it follows that Maria is aware of nothing in having her B experience and aware of nothing in having her Y experience.

Now we can see the first worry raised by the Maria case. Take her first experience. I hold that Maria has a B visual experience even though she is visually aware of nothing in having her experience. (She is *visually* aware of nothing. But she might be concurrently, say, aurally aware of various sounds and sound properties.) This is the first bizarre consequence: my form of pure phenomenism entails that *there can be something it’s like for Maria at a time t even though she is aware of nothing at t .*¹¹⁷

¹¹⁷ Related to this consequence are two other surprising consequences:

First, my view entails that there are *four*, not three, distinct types of perceptual experience. Most philosophers accept the following taxonomy of experiences. Best of all are veridical experiences, in which you succeed in perceiving an object and the object has all of the properties it seems to have. Worse are illusions, in which you succeed in perceiving an object but it lacks some of the properties it seems to have. Worst of all are hallucinations, in which you fail to perceive any object at all.

My view entails a fourth sort of experience: call them “extreme hallucinations.” Maria is having an

Philosophers and non-philosophers alike tend to use the words ‘consciousness’ and ‘awareness’ to mean the same thing. To the extent that phenomenal consciousness is viewed simply as a form of awareness, then, my position’s consequence—that you can be phenomenally conscious even though you are aware of nothing—will strike many as contradictory. How, they might ask, could you be both conscious (read: aware of something) and not aware of anything? Surely this is a *reductio* of my position.

If *this* is the reason for thinking that the counterintuitive consequence is a reason for giving up my position, I don’t think it is a very good one. When we use ‘conscious’ and ‘aware’ interchangeably, we are speaking loosely.¹¹⁸ Though consciousness might often be *accompanied* by awareness, it does not follow that it *is* awareness.

Moreover, I think that our tendency to use ‘conscious’ and ‘aware’ as synonyms reflects the pull that the Confrontation Model has on our theoretical judgments about consciousness. Remember: there is something about the very phenomenology of experience, I claim, that makes the Confrontation Model irresistible. Since the Model makes brute confrontation—a form of awareness—necessary for phenomenality, it is no wonder that, given the Model’s perennial attractiveness, we are liable to regard consciousness as just a kind of awareness. So I do not think that this consequence of my view is problematic.

extreme hallucination in having her B (and Y) experience. She not only fails to see any object, but she also fails to represent a property (blue), though she is undergoing an experience with a phenomenal property, B. Look at it this way. In a hallucination, your experience represents your environment as instantiating properties which it fails to instantiate. In an extreme hallucination, though, you don’t even succeed in doing this much, that is, misrepresenting your environment. You have a representationally “blank” but phenomenally conscious experience.

Second, you might have thought that, though you can’t hallucinate Tyler unless you have seen him, you *can* hallucinate green even if you have never seen green things. If I’m right, though, and pure phenomenism, phenomenal internalism, and strong content externalism are true, then you can’t hallucinate green unless you have seen green things. You could have an experience phenomenally indistinguishable from a hallucination of green, but that isn’t the same. Maria’s B experience, then, is *not* a hallucination of blue. It is something else entirely: an extreme hallucination.

¹¹⁸ Of course, people are free to go on saying that phenomenal consciousness is awareness if they like, provided that by ‘awareness’ they mean neither awareness of objects nor properties.

The Second Counterintuitive Consequence

But there is another worry. How, it will be asked, could Maria possibly notice the phenomenal change from t_1 to t_2 if, in having her experiences, she is aware of nothing? This is the second bizarre consequence: my view entails that *Maria can notice the phenomenal change in the course of her visual experience even though she is aware of nothing in having her experiences before and after the change.*

In slightly more detail, my view is inconsistent with the following principle—call it the “Change Principle”:

CHANGE PRINCIPLE—A subject can notice a phenomenal change in the course of her experience within a particular perceptual modality only if either before or after the change she is aware of something in having her experience(s).

Notice two things about the Change Principle.

First, the Change Principle is not the stronger principle that a subject can notice a phenomenal change only if she is aware of something *before* the change. Suppose you, not Maria, are in a deep, dreamless sleep. (You can allow in dreams as long as you don’t think dreams are experiences.) While sleeping, you are visually aware of nothing. Upon being jolted from your sleep by a sudden loud noise, though, you are visually aware of something. What is important, here, is that you notice this phenomenal change even though, before, you were visually aware of nothing. How? By noticing the sudden visual *presence* of a range of objects and properties.¹¹⁹

Second, the Change Principle is not the stronger principle that a subject can notice a phenomenal change only if she is aware of something *after* the change. Suppose you, not Maria, are in the neurosurgeon’s office staring up at the ceiling as he works on your brain in an attempt

¹¹⁹ I owe the waking up example to Steve Yablo.

to cure the constant tingle in your hands. While you admire the white ceiling, he accidentally puts you in a brain state that causes you to have the visual experience you would have if you were to see something vermilion, a shade of red you have had no causal contact with, let's suppose.¹²⁰ This apparent shade fills your visual field. Before the surgeon's slip, you were visually aware of the ceiling and the color white. After it, you are visually aware of nothing. But surely you notice the change. How? By noticing the sudden visual *absence* of the ceiling and white.

So the Change Principle is fairly weak.

My collection of views entail that Maria is visually aware of nothing before and after the change. By the Change Principle, she ought not to notice the phenomenal change. But she does, so, given the Change Principle, something in my collection of views has to go.

The problem here is simple. Pure phenomenism, phenomenal internalism, strong content externalism, and the Awareness Principle are incompatible with the Change Principle. Representationists accept the Awareness Principle and the Change Principle. They give up pure phenomenism and, depending on whether their view is narrow or broad, one or the other of phenomenal internalism or strong content externalism. They would insist that the only way to salvage pure phenomenism in the face of the Maria case is for pure phenomenists to give up the Awareness Principle and say that Maria is aware of properties in having her experiences. Since pure phenomenists are not sense-datum theorists, those properties could only be properties of her experience—her own phenomenal properties. So representationists would say that the Maria case shows that pure phenomenists must hold a view that conflicts with transparency after all.

I think that pure phenomenists can avoid this by rejecting, not the Awareness Principle, but the Change Principle. But if I want to give up the change principle, I have to show how

¹²⁰ Let's assume that you are unable to gain vermilion content by interpolation.

Maria can notice the phenomenal change even though neither before nor after the change she is visually aware of anything in having her B and Y experiences.

There are two ways to do it.

First, the Change Principle rules out brute awareness of change, awareness of change not grounded in awareness of things that change. But why rule this out? Since brute awareness of change certainly seems plausible enough, the Change Principle is unmotivated. But if it is unmotivated, we have been given no good reason for thinking that the second counterintuitive consequence is a reason for rejecting my position.

Second, even if we grant that brute awareness of change is impossible, the Change Principle still has a big problem. This will take a bit of explaining.

‘Concentrating intently in having her experiences’ is ambiguous. It could mean concentrating intently on what Maria might call the scene before her eyes. Call this “outward-concentrating.” Or it could mean concentrating intently on her experience itself. This is just introspecting. We could also call it “inward-concentrating.”

I will proceed as follows. First, I will explain what I think introspection is. The notion of introspection will play an important role in my account of how Maria notices the phenomenal change. Second, I’ll suppose that Maria is inward-concentrating—introspecting—during the phenomenal change and explain how she could notice it. Finally, I’ll suppose that she is outward-concentrating during the phenomenal change and explain how she could notice it.

Introspection

I believe that whenever a subject undergoes a perceptual experience within a particular sense modality he enters into a subpersonal state with representational content about his

perceptual experience. Take Jones. During a certain interval of time, Jones has three visual experiences in quick succession: V_1 , V_2 , then V_3 . At the same time, Jones enters in quick succession three subpersonal states with content about these experiences: S^V_1 , S^V_2 , then S^V_3 . S^V_1 's content is that V_1 has such and such phenomenal properties; S^V_2 's content is that V_2 has such and such phenomenal properties; and S^V_3 's content is that V_3 has such and such phenomenal properties. Call these subpersonal states "shadowing states" because of the way they accompany the perceptual experiences they are about: they are the doxastic shadows of phenomenally conscious perceptual experiences. (Actually, I think that bodily sensations, moods, and other phenomenally conscious mental events—not just perceptual experiences—have shadowing states.)

Subjects can form occurrent thoughts about their experiences, the contents of which come from the experiences' corresponding shadowing states. Jones, for example, can form an occurrent thought T_1 about his visual experience V_1 that takes V_1 's shadowing state S^V_1 's content as its own. T_1 's content is that V_1 has such and such phenomenal properties.

Very important: to have one of these inner-directed occurrent thoughts a subject needn't have the concept of phenomenal property or be able to say in words things like 'My visual experience has the phenomenal property B'.

Now I can say what introspection is. Jones introspects his visual experience V_1 if he entertains a V_1 -directed occurrent thought T_1 at the same time as he has V_1 . If Jones thinks T_1 while he's undergoing V_1 , he counts, on my view, as introspecting V_1 . Thus, I think that introspecting is a kind of thinking about one's experiences, thinking that happens concurrently with having the experiences one thinks about.

One more thing. Subjects can form occurrent thoughts about *past* perceptual experiences,

thoughts which take for their contents the contents of the past experiences' corresponding shadowing states. As Jones undergoes V_3 , he introspects it: that is, he thinks the occurrent thought T_3 , which takes its content from S^V_3 . But, even as he does this, he can think a thought about, say, V_2 , the experience that immediately preceded V_3 . This thought— T_M —takes S^V_2 's content. By entertaining both T_3 and T_M , Jones can introspectively compare his current visual experience with the one that preceded it, coming to know, for example, that V_3 differs phenomenally in various ways from V_2 . Call thoughts like T_M "introspective memories."¹²¹

Explaining Maria's Predicament

With this little theory of introspection in hand, I can proceed to explain how Maria could notice the phenomenal change without being aware of anything in having her experiences.¹²²

Suppose Maria is inward-concentrating during the phenomenal change: both before and after the switch she is introspecting her experiences. Before the change, in introspecting her B experience, she thinks an occurrent thought T_B with the content that her experience is B. After the change, in introspecting her Y experience, she thinks an occurrent T_Y with the content that her experience is Y. But, since she was concentrating throughout, Maria is able to detect the difference in content between T_B and T_Y . In virtue of noticing this difference, she notices the

¹²¹ Introspective memories aren't always easy to have. Our introspective memories go back only so far. And sometimes we can't even remember very recent past experiences.

Carruthers (1989) considers such a case. While driving home one evening, Carruthers is so caught up in thoughts about work that his "conscious attention" is "wholly abstracted from [his] surroundings" (p.258). "It is common in such cases," he observes, "that one may suddenly 'come to,' returning one's attention to the task at hand with a startled realization that one has not the faintest idea what one has been doing or seeing for some minutes past" (ibid).

I'm not sure what the difference is between cases like Jones's, where subjects are able to access recent introspective memory contents, and Carruthers cases, where subjects aren't able to do this. I suspect it has something to do with the fact that subjects in the former cases are attentive to either their experiences or their environments in a way that subjects in the latter cases aren't. The Carruthers cases involve subjects who are somehow absentminded. See below.

¹²² I recognize that I have offered, at best, only a proto-theory of introspection. I plan to elaborate my views on introspection in future work.

phenomenal difference between her two experiences.

(In the nondeviant case—the case where Maria’s experiences have b and y content, respectively, and she is aware of these properties in having them—the noticing is *overdetermined*. Both (1) the difference between T_B and T_Y and (2) the difference between pre-change and post-change experiential content cause the noticing.)

Suppose Maria is outward-concentrating during the phenomenal change: before and after the switch she is concentrating intently on what she thinks of as the scene before her eyes. In this case, though she does not begin by introspecting, the phenomenal change is so, well, big that it intrudes on her psychology in an introspection-prompting way: there is such a phenomenal jolt that she immediately entertains the occurrent thought T_Y about her Y experience.¹²³ Now, even though she was not introspecting before the phenomenal change, she was nevertheless *concentrating*, engaged in a kind of intense mental attention. Because of this, once she begins introspecting (entertaining T_Y while having her Y experience), she is able to access introspective memories.¹²⁴ Among these introspective memories is the occurrent thought T_B . In virtue of noticing the difference between T_B and T_Y , Maria notices the phenomenal change between her two experiences.

(Again, in the non-deviant case the noticing is *overdetermined*. Both (1) the difference between T_B and T_Y and (2) the difference between pre-change and post-change experiential

¹²³ The phenomenal change might not have been so big; there might have been no jolt. Tiny phenomenal changes could very easily escape notice.

What do I mean by a “big change” or a “jolt”? Consider two courses of experience A and B. In A, you go from having your entire visual field filled with blue to having it instantaneously filled with yellow. In B, you go from having an experience as of a forest to having an experience as of the same forest, unchanged but for a tiny movement in the grass down towards the bottom right of your visual field. There is an intuitive sense in which A involves a bigger phenomenal change than B. It’s much easier to notice the change in A than it is to notice the change in B.

¹²⁴ Why does concentration, whether inner- or outer-directed, help make introspective memories available? I don’t know—though that it does is extremely plausible.

content cause the noticing.)

Thus, even if Maria is unaware of anything in having her experience, she can still notice the phenomenal change. She can do this by noticing an introspective difference. The key here is that, since pure phenomenism makes the phenomenal properties B and Y intrinsic properties, Maria is guaranteed, even by content externalism's lights, to be able to have mental events and states with B and Y content—even with her extremely limited causal history.

Think of it this way. The properties Maria is aware of *in having her experiences* are the properties, if any, her experiences represent her environment as instantiating. Since her experiences have no content, she's aware of no properties in having her experiences. But the occurrent thoughts Maria entertains when she introspects also involve awareness. *In introspecting*—that is, in entertaining occurrent thoughts about her experiences while undergoing them—Maria is aware of properties: her experiences' phenomenal properties.

Now recall the Change Principle:

CHANGE PRINCIPLE—A subject can notice a phenomenal change in the course of her experience within a particular perceptual modality only if either before or after the change she is aware of something in having her experience(s).

Note the wording. Friends of the Change Principle assume that if a subject notices a phenomenal change between experiences in a given sense modality, then the awareness necessary for the noticing is the sort constitutive of experiences in *that* modality. But why assume this? That is, why assume that Maria can notice a phenomenal difference between two of her visual experiences only if she is aware of something *in having those visual experiences*? I claim that this assumption is unfounded. Maria does not notice the phenomenal change by being *visually* aware of things that change. She notices it by being *introspectively* aware of things that change.

Now let me reply to an anticipated objection. You might be thinking something like the

following: “Your view conflicts with transparency after all. Don’t these remarks about introspection contradict the transparency-based Awareness Principle, the principle that all subjects are aware of in having their experiences are the objects (if any) they perceive and the properties (if any) their experiences represent those objects as having?”

No. Again, the transparency point—the one that motivates the Awareness Principle—is simply that when you introspect your experiences all you will find you are aware of *in having your experiences* are the objects (if any) you perceive and the properties (if any) your experiences represent. The ‘in having your experiences’ is crucial, for the mental act of awareness involved in experiencing is a different mental act from the mental act of awareness involved in introspecting an experience. In having your G experience of the avocado, you are spontaneously inclined to judge that a property, g, is presented and located. But introspection is a different type of mental act, one, moreover, *without experiential phenomenology*. In introspecting an experience, you are not spontaneously inclined to judge that any properties are presented and located. You entertain content about your experience, but in doing so the phenomenal lights, so to speak, are not on.¹²⁵

Consider again this passage from Tye:

Your experience is thus transparent to you. When you try to focus upon it, you ‘see’ right through it, as it were, to the things apparently outside and their apparent qualities (2002, p.139).

This comes after Tye makes what I have been calling the transparency point. Note, though, that what he says here does not entail that introspection is not a kind of representational awareness of experiencing. That would follow only if, upon introspecting your experience E of the avocado, you found that things seemed different to you and that either old properties now seemed to

¹²⁵ There is *some* phenomenology associated with introspecting. I think that this phenomenology comes, not from introspecting, but from the experiences of “inner saying” that usually—if not always—accompany introspection. See chapter 2.

belong to your experience or new properties seemingly belonging to your experience popped into view. But of course that isn't what it's like! The reason why—and the reason its not being like that is consistent with regarding introspection as awareness of properties of experience—is that introspection does not have perceptual phenomenology. The reason why what it's like does not change when you introspect and things continue visually unchanged for you when you begin introspecting E is that introspective states do not have a perceptual “what it's like” that could alter or somehow intrude upon your perceptual phenomenology.

So Maria can notice the phenomenal change from her having a B experience at t_1 to her having a Y experience at t_2 , even though she is aware of nothing in having each of those experiences, because she is capable of introspecting her experiences. She notices the difference, in short, because she notices differences in what she is introspectively aware of.

6. Conclusion of the Dissertation

I have presented two competing models of phenomenality, the Confrontation Model and the Possession Model, and argued that we ought to accept the latter. I have defended a particular application of the Possession Model, pure phenomenism, and shown, not only that it does not conflict with the transparency of experience, but that it is compatible with direct realism. I conclude that we should hold pure phenomenism. Phenomenality cannot be explained in terms of intentionality.

This dissertation left several loose ends. Chief among these were my criticisms of impure representationism and impure phenomenism. In future research, I hope to shore up my case *for* phenomenal internalism and strong content externalism and my case *against* impure representationism and impure phenomenism.

References

- Alston, William. (1998) "Perception and Conception." In *Pragmatism, Reason, and Norms*, ed. Kenneth R. Westphal. New York: Fordham University Press.
- _____. (1999) "Back to the Theory of Appearing." In *Philosophical Perspectives* 13, ed. James E. Tomberlin. Cambridge: Blackwell.
- Anscombe, G.E.M. (1965) "The Intentionality of Sensation: A Grammatical Feature." In *Analytical Philosophy, Second Series*, ed. R.J. Butler. Oxford: Basil Blackwell.
- Aristotle. (1984) *De Anima*. In *The Complete Works of Aristotle: The Revised Oxford Translation*, ed. Jonathan Barnes. Princeton: Princeton University Press.
- Armstrong, David. (1968) *A Materialist Theory of the Mind*. London: Routledge & Kegan Paul.
- Arnauld, Antoine. (1990) *On True and False Ideas*. Trans. Elmar J. Kremer. Lewiston/Queenston/Lampeter: Edwin Mellen Press.
- Austin, John. (1962) *Sense and Sensibilia*. Oxford: Oxford University Press.
- Baldwin, Thomas. (1990) *G.E. Moore*. London: Routledge.
- Barnes, W.H.F. (1944-45) "The Myth of Sense-Data." *Proceedings of the Aristotelian Society* 45. Reprinted in Swartz (1965).
- Bennett, Jonathan. (1971) *Locke, Berkeley, Hume*. Oxford: Clarendon.
- Block, Ned. (1990) "Inverted Earth." In *Philosophical Perspectives* 4, ed. James E. Tomberlin. Atascadero: Ridgeview. Reprinted in Block, Flanagan, and Güzeldere (1997).
- _____. (1995) "On a Confusion About a Function of Consciousness." *Behavioral and Brain Sciences* 18. Reprinted in Block, Flanagan, and Güzeldere (1997).
- _____. (Forthcoming) "Mental Paint." In *Others on Burge: 10 Essays With Responses From Tyler Burge*, ed. Martin Hahn and Bjørn Ramberg. Cambridge, MA: MIT Press.
- Block, Ned, Owen Flanagan, and Güven Güzeldere (eds.). (1997) *The Nature of Consciousness*. Cambridge, MA: MIT Press.

- Broad, C.D. (1923) *Scientific Thought*. London: Routledge and Kegan Paul.
- Burge, Tyler. (1979) "Individualism and the Mental." In *Midwest Studies in Philosophy* 4, eds. P. French, T. Uehling, and H. Wettstein. Minneapolis: University of Minnesota Press.
- Byrne, Alex. (2001) "Intentionalism Defended." *Philosophical Review* 110.
- _____. (2003) "Consciousness and Nonconceptual Content." *Philosophical Studies* 113.
- _____. (2004) "What Phenomenal Consciousness Is Like." In *Higher-Order Theories of Consciousness*, ed. R. Gennaro. Amsterdam/Philadelphia: John Benjamins Publishing Company.
- Byrne, Alex, and David R. Hilbert. (1997) *Readings on Color, Volume 1: The Philosophy of Color*. Cambridge, MA: MIT Press.
- Carruthers, Peter. (1989) "Brute Experience." *Journal of Philosophy* 86.
- Chalmers, David. (1996) *The Conscious Mind*. Oxford: Oxford University Press.
- _____. (2002) "The Components of Content." In *Philosophy of Mind: Classical and Contemporary Readings*, ed. David Chalmers. Oxford: Oxford University Press.
- _____. (2004a) "The Representational Character of Experience." In *The Future for Philosophy*, ed. B. Leiter. Oxford: Oxford University Press.
- _____. (2004b) "Perception and the Fall from Eden." MS.
- Child, William. (1992) "Vision and Experience: The Causal Theory and the Disjunctive Conception." *Philosophical Quarterly* 42.
- Chisholm, Roderick. (1957) *Perceiving: A Philosophical Study*. Ithaca, NY: Cornell University Press.
- Crane, Tim. (1992) (ed.) *The Contents of Experience*. Cambridge: Cambridge University Press.
- Crane, Tim, and D.H. Mellor. (1990) "There is No Question of Physicalism." *Mind* 94.
- Dennett, Daniel. (1988) "Quining Qualia." In *Consciousness in Contemporary Science*, ed. A.J. Marcel and E. Bisiach. Oxford: Oxford University Press. Reprinted in Block, Flanagan, and Güzeldere (1997).

- Dretske, Fred. (1981) *Knowledge and the Flow of Information*. Oxford: Basil Blackwell.
- _____. (1995a) *Naturalizing the Mind*. Cambridge, MA: MIT Press.
- _____. (1995b) "Phenomenal Externalism." In *Philosophical Issues 7*, ed. Enrique Villanueva. Atascadero: Ridgeview.
- _____. (1999) "The Mind's Awareness of Itself." *Philosophical Studies* 95.
- Ducasse, C.J. (1951) *Nature, Mind, and Death*. Illinois: Open Court.
- Egan, Andy. (2004) "Appearance Properties?" MS.
- Egan, Andy, and James John. (2004) "A Puzzle about Perception." MS.
- Evans, Gareth. (1980) *The Varieties of Reference*. Oxford: Oxford University Press.
- Fodor, Jerry. (1975) *The Language of Thought*. Cambridge, MA: Harvard University Press.
- Foster, John. (2000) *The Nature of Perception*. Oxford: Clarendon.
- _____. (1987) *Psychosemantics*. Cambridge, MA: MIT Press.
- Grice, H.P. (1961) "The Causal Theory of Perception." *Proceedings of the Aristotelian Society*, Supp. Vol. 35.
- Hammond, Michael, Jane Howarth, and Russell Keat. (1991) *Understanding Phenomenology*. Oxford: Basil Blackwell.
- Harman, Gilbert. (1990) "The Intrinsic Quality of Experience." In *Philosophical Perspectives* 4, ed. James E. Tomberlin. Atascadero: Ridgeview. Reprinted in Block, Flanagan, and Güzeldere (1997).
- _____. (1996) "Explaining Color in Terms of Subjective Reactions." In *Philosophical Issues 7*, ed. Enrique Villanueva. Atascadero: Ridgeview. Reprinted in Byrne and Hilbert (1997).
- Heck, Richard. (2000) "Nonconceptual Content and the Space of Reasons." *Philosophical Review* 109.
- Hilbert, David R., and Mark E. Kalderon. (2000) "Color and the Inverted Spectrum." In *Color Perception*, ed. S. Davis. Oxford: Oxford University Press.
- Hinton, J.M. (1973) *Experiences*. Oxford: Clarendon.

- Hume, David. (1977) *An Enquiry Concerning Human Understanding*. Indianapolis: Hackett Publishing Company.
- Jackson, Frank. (1977) *Perception: A Representative Theory*. Cambridge: Cambridge University Press.
- John, James. (2004) "Indirect Realism and the Veil of Perception." MS.
- Johnston, Mark. (2004) "The Obscure Object of Hallucination." *Philosophical Studies* 120.
- Langsam, Harold. (1995) "The Theory of Appearing Defended." *Philosophical Studies* 87.
- Levine, Joseph. (2001) *Purple Haze*. Oxford: Oxford University Press.
- Lewis, David. (1994) "Reduction of Mind." Reprinted in Lewis (1999).
- _____. (1999) *Papers in Metaphysics and Epistemology*. Cambridge: Cambridge University Press.
- Locke, John. (1964) *An Essay Concerning Human Understanding*. Ed. A.D. Woozley. London.
- Lycan, William. (1996) *Consciousness and Experience*. Cambridge, MA: MIT Press.
- _____. (2001) "The Case for Phenomenal Externalism." In *Philosophical Perspectives* 15, ed. James E. Tomberlin. Boston: Blackwell.
- Mackie, John. (1976) *Problems from Locke*. Oxford: Clarendon.
- Martin, M.G.F. (1997) "The Reality of Appearances." In *Thought and Ontology*, ed. Mark Sainsbury. Milan: FrancoAngeli.
- _____. (2002) "The Transparency of Experience." *Mind & Language* 17.
- _____. (2004) *Uncovering Appearances*. MS.
- Moore, G.E. (1993) "The Refutation of Idealism." In *G.E. Moore: Selected Writings*, ed. Thomas Baldwin. London: Routledge.

- McDowell, John. (1982) "Criteria, Defeasibility, and Knowledge." Reprinted in McDowell (1998).
- _____. (1986) "Singular Thought and the Extent of Inner Space." Reprinted in McDowell (1998).
- _____. (1994) *Mind and World*. Cambridge, MA: Harvard University Press.
- _____. (1998) *Meaning, Knowledge, and Reality*. Cambridge, MA: Harvard University Press.
- McGinn, Colin. (1989) "Can We Solve the Mind-Body Problem?" *Mind* 98. Reprinted in Block, Flanagan, and Güzeldere (1997).
- _____. (1991) *The Problem of Consciousness*. Oxford: Blackwell. "Consciousness and Content" reprinted in Block, Flanagan, and Güzeldere (1997).
- _____. (1997) "Critical Study of Dretske's *Naturalizing the Mind*." *Noûs* 31.
- Nagel, Thomas. (1974) "What Is It Like to Be a Bat?" *Philosophical Review* 83.
- _____. (1986) *The View from Nowhere*. Oxford: Oxford University Press.
- Peacocke, Christopher. (1983) *Sense and Content*. Oxford: Oxford University Press.
- _____. (1992) *A Study of Concepts*. Cambridge, MA: MIT Press.
- Pitcher, George. (1970) *A Theory of Perception*. Princeton: Princeton University Press.
- Price, H.H. (1932) *Perception*. London: Methuen.
- Putnam, Hilary. (1975) "The Meaning of 'Meaning'." In *Minnesota Studies in the Philosophy of Science* 2, ed. Keith Gunderson. Minneapolis: University of Minnesota Press.
- Rey, Georges. (1998) "A Narrow Representationalist Account of Qualitative Experience." In *Philosophical Perspectives* 12, ed. James E. Tomberlin. Oxford: Blackwell.
- Robinson, Howard. (1994) *Perception*. London: Routledge.
- Russell, Bertrand. (1912) *The Problems of Philosophy*. Indianapolis: Hackett Publishing Company.

- Ryle, Gilbert. (1949) *The Concept of Mind*. London: Hutchinson.
- Searle, John. (1983) *Intentionality*. Cambridge: Cambridge University Press.
- Sellars, Wilfred. (1956) "Empiricism and the Philosophy of Mind." In *Minnesota Studies in the Philosophy of Science* 1, ed. Herbert Feigl and Michael Scriven. Minneapolis: University of Minnesota Press.
- Shoemaker, Sydney. (1982) "The Inverted Spectrum." *Journal of Philosophy* 74.
Reprinted in Block, Flanagan, and Güzeldere (1997).
- _____. (1994) "Phenomenal Character." *Noûs* 28.
- _____. (2002) "Content, Character, and Color." Whitehead Lectures, Harvard University.
- Siewert, Charles. (1998) *The Significance of Consciousness*. Princeton: Princeton University Press.
- Smart, J.J.C. (1959) "Sensations and Brain Processes." *Philosophical Review* 68.
- Smith, A.D. (2002) *The Problem of Perception*. Cambridge, MA: Harvard University Press.
- Snowdon, Paul F. (1980-81) "Perception, Vision, and Causation." *Proceedings of the Aristotelian Society* 81.
- Stalnaker, Robert. (1984) *Inquiry*. Cambridge, MA: MIT Press.
- _____. (1998) "What Might Nonconceptual Content Be?" In *Philosophical Issues* 9, ed. Enrique Villanueva. Atascadero: Ridgeview.
- _____. (1999) "Comparing Qualia Across Persons." Reprinted in Stalnaker (2003).
- _____. (2003) *Ways a World Might Be: Metaphysical and Anti-Metaphysical Essays*. Oxford: Oxford University Press.
- Stevens, Wallace. (1923) *Harmonium*. New York: Alfred A. Knopf.
- Stich, Stephen and Ted Warfield (eds.). (1994). *Mental Representation: A Reader*. Oxford: Blackwell.
- Stoljar, Daniel. (Forthcoming) "The Argument from Diaphanousness." In *Language, Mind, and World: Special Issue of the Canadian Journal of Philosophy*, ed. Maite Ecurdia, Robert Stainton, and Chris Viger.

- Swartz, Robert J. (1965) (ed.) *Perceiving, Sensing, and Knowing*. Los Angeles: University of California Press.
- Thau, Michael. (2002) *Consciousness and Cognition*. Oxford: Oxford University Press.
- Tye, Michael. (1995) *Ten Problems of Consciousness*. Cambridge, MA: MIT Press.
- _____. (2000) *Consciousness, Color, and Content*. Cambridge, MA: MIT Press.
- _____. (2002) "Representationalism and the Transparency of Experience." *Noûs* 36.
- Valberg, J.J. (1992) *The Puzzle of Experience*. Oxford: Oxford University Press.
- Williamson, Timothy. (2000) *Knowledge and its Limits*. Oxford: Oxford University Press.
- Wilson, Margaret. (1999) *Ideas and Mechanism*. Princeton: Princeton University Press.
- Yolton, John. (1984) *Perceptual Acquaintance from Descartes to Reid*. Minneapolis: University of Minnesota Press.
- _____. (1996) *Perception and Reality: A History from Descartes to Kant*. Ithaca, NY: Cornell University Press.