UNIVERSITY OF CALIFORNIA, SAN DIEGO

Sensory Underdetermination and Perceptual Constancy

A Dissertation submitted in partial satisfaction of the
requirements for the degree
Doctor of Philosophy

in

Philosophy and Cognitive Science

by

Damon Crockett

Committee in charge:

Professor Jonathan Cohen, Chair
Professor William Bechtel
Professor Matthew Fulkerson
Professor Donald MacLeod
Professor John Serences

2015

The Dissertation of Damon Crockett is approved, and it is acceptable in quality and form for publication on microfilm and electronically:

_____

_____

_____

_____

_____

Chair

University of California, San Diego

2015

TABLE OF CONTENTS

Signature Page iii

Table of Contents iv

List of Figures vi

Acknowledgements vii

Vita ix

Abstract of the Dissertation x

Chapter 1 Phenomenal Content 1
- 1.1 The Phenomenology of Perceptual Experience 1
- 1.2 The Representational Contents of Perceptual Experience . 3
- 1.3 Intentionalism 5
- 1.4 Phenomenal Content 7
- 1.5 Explanation 10
  - 1.5.1 Supervenience Is Too Strong 10
  - 1.5.2 Supervenience Is Too Weak 16
- 1.6 Perception and Cognition 18

Chapter 2 Indeterminacy 21
- 2.1 Clarification of the Thesis 23
- 2.2 Phenomenal Content Cannot be Indeterminate 26
  - 2.2.1 The *L*-Supervenience of Phenomenology on Content is a Condition on Explanation 28
  - 2.2.2 Phenomenology Cannot Explain Indeterminate Contents 29
  - 2.2.3 Indeterminate Contents Cannot Explain Phenomenology 31
  - 2.2.4 The *L*-Supervenience of Content on Phenomenology is a Condition on Explanation 40
- 2.3 Siegel and the Method of Phenomenal Difference 43
  - 2.3.1 The Method 43
  - 2.3.2 Why Not Ducks? 44
- 2.4 Morrison and Perceptual Confidence 50
  - 2.4.1 Perceptual Confidence 51
  - 2.4.2 Perceptual Confidence and Doxastic Confidence . . 53
  - 2.4.3 Assessing Morrison's Argument 54
- 2.5 Wrap-up 63

Chapter 3 Underdetermination . . . 64
3.1 Constancy and Phenomenal Content . . . 64
3.2 Sensory Underdetermination . . . 65
3.2.1 The Problem in General . . . 65
3.2.2 Stacking Problems . . . 68
3.3 Phenomenalism . . . 70
3.4 Phenomenal Content and Overreach . . . 74
3.4.1 Sensory and Phenomenal . . . 74
3.4.2 Grouping and Depth (and Ducks Again) . . . 76
3.4.3 Energy Map and Worldmaker Contents . . . 81
3.4.4 The Role of Feedback in Awareness . . . 85
3.4.5 Two Types of Overreach . . . 89
3.5 Looking Ahead . . . 92

Chapter 4 Constancy . . . 93
4.1 A Phenomenal Puzzle . . . 94
4.2 The Standard View: Full Phenomenal Resolution . . . 97
4.3 Against SV . . . 101
4.3.1 The Phenomenal Trade-Off . . . 101
4.3.2 Phenomenal Signatures . . . 107
4.3.3 Matching Behavior . . . 112
4.4 Indeterminate Content . . . 120
4.4.1 Indeterminacy and Symmetric Comparisons . . . 122
4.4.2 Determinacy Within, Indeterminacy Across . . . 124
4.4.3 Indeterminacy Everywhere Redux . . . 128
4.5 Final Remarks . . . 129

Chapter 5 Denouement . . . 131
5.1 Partial Phenomenal Resolution . . . 132
5.1.1 The Puzzle Intuitions . . . 132
5.1.2 Overreach in the Energy Map . . . 136
5.1.3 Worldmaker Resolution . . . 142
5.2 The Puzzle Intuitions Without Phenomenal Resolution . . . 158
5.3 Summary and Concluding Remarks . . . 169

Bibliography . . . 171

LIST OF FIGURES

Figure 1.1: The Mueller-Lyer line illusion. 19

Figure 2.1: Stucco walls. We experience these differently because we represent them as *being* different. 27
Figure 2.2: Fork. In this structure, c does not determine either e. 28
Figure 2.3: A 'color curve', which represents a single color content on the hypothetical 'confidence' theory of color. 37

Figure 3.1: An example of ambiguous grouping. 77
Figure 3.2: A possible schematic for two processing pathways. 83
Figure 3.3: A Necker cube. 83
Figure 3.4: Kanizsa-type figure. 86
Figure 3.5: The Gelb effect. 91

Figure 5.1: The Adelson checkerboard illusion. 137
Figure 5.2: Lifter. The picture is ambiguous between two grouping interpretations, one on which the hanging arms belong to the lifter and another on which they belong to the spotter behind him. 147

ACKNOWLEDGEMENTS

First and foremost, I'd like to thank Jonathan Cohen, my advisor and committee chair, without whose encouragement and guidance a doctoral degree would never have been possible. His work in the philosophy of perception was both a model for good philosophy and a catalyst for much of the discussion herein. He taught me to be undogmatic and to focus on arguments instead of labels or approaches.

I'd like also to thank the members of the Philosophy of Perception Reading Group at UCSD for reading most of this document—in a very early and quite terrible form—and giving useful feedback. I'd like to mention Matthew Fulkerson in particular, who joined my committee late in the game but has been a big help nonetheless, both by giving challenging feedback and by expanding my view of the discipline.

I'd like to give special thanks to the Vision Journal Club at UCSD, and in particular to Don MacLeod, John Serences, and Karen Dobkins. Everything I know about visual processing I owe to my time with that group, and in fact, some of the larger bits of argumentative structure herein starting taking form as a result of our discussions. The VJC embodies the interdisciplinary spirit of UCSD, and I'm very grateful that philosophers are allowed to join the discussion, ignorant as we are.

I'd like also to thank Bill Bechtel, whose teaching and work on explanation in science had a big influence on my thinking here.

I'd like to thank Karen Underwood for her patience and love during a time when I was a reclusive and cantankerous partner.

I'd like to thank my grandmother, Gretchen Hochberger, for her boosterism.

And I'd like to thank my parents, Peter and Lori Crockett, for all their support and understanding.

Portions of chapter 4 have significant overlap with my article, published in the *Australasian Journal of Philosophy*, entitled "Surface Colour is Not a Perceptual Content" [Crockett, 2015].

VITA

2007	B. A. in Philosophy, *High Distinction*, University of Illinois at Urbana-Champaign
2015	Ph. D. in Philosophy and Cognitive Science, University of California, San Diego

PUBLICATIONS

Damon Crockett, "Surface Colour is Not a Perceptual Content", *Australasian Journal of Philosophy*, 93/2: 303-318, 2015.

ABSTRACT OF THE DISSERTATION

Sensory Underdetermination and Perceptual Constancy

by

Damon Crockett

Doctor of Philosophy in Philosophy and Cognitive Science

University of California, San Diego, 2015

Professor Jonathan Cohen, Chair

This project has as its focus a pair of related phenomena central to human perception. The first is the underdetermination of perceptual content by sensor input, and the second is a class of mechanisms designed to transform impoverished sensor input into useful perceptual content, mechanisms commonly called 'perceptual constancies'. The goal of this project is to discuss a particularly difficult form of sensory underdetermination I call *stacking*, a *co-local* sensory conflation of distal properties like surface color and illumination, or size and distance. And although stacking problems are not computationally intractable, there appear to be significant constraints on

potential solutions to these problems, constraints rooted in the phenomenal structure of perceptual experience. Accordingly, I spend much of the project examining in detail a type of intentional content I argue explains the phenomenology of perceptual experience—*phenomenal content*. With explanatory adequacy as a guiding principle, I argue that phenomenal contents cannot be indeterminate, in the sense that they cannot fail to be attributional. I also argue that phenomenal contents must be *phenomenally bounded*, a conclusion that rules out of consideration various high-level properties like natural kinds. Additionally, I argue that there are in fact two distinct domains of phenomenal content. The first I call the *energy map*, as its contents are tied closely to the distribution of energy across the sensors. The second I call the *worldmaker*, as its contents are what make the world accessible and intelligible to perceivers. Worldmaker contents are underdetermined by energy map contents and are a phenomenal manifestation of perceptual mechanisms of *overreach*, the general term for solutions to problems of sensory underdetermination. After setting up these constraints, I apply them to the problem of color constancy, and I argue that the standard view of color constancy, found both in philosophy and in perceptual psychology—the view that we perceptually represent spectral surface reflectance—must be wrong. My methodology and conclusions have implications both for philosophical accounts of perception and for computational models of perception.

Chapter 1

Phenomenal Content

The central purpose of this chapter is to discuss *phenomenal perceptual content*, which notion will be the focus of this project. Of particular importance for the project is the role that phenomenal perceptual content might play in explaining the phenomenology of perceptual experience. But before we can discuss that role, we need to know what phenomenal perceptual content is. And before we can do *that*, we need first to introduce some other of our terms.

1.1 The Phenomenology of Perceptual Experience

The term 'phenomenal', as it is used here, is connected with the term 'phenomenology', which has technical meaning in philosophy. Confusingly, 'phenomenology' picks out both a particular philosophical approach and a feature of our mental lives. Our focus will be on the latter of these, and I will henceforth have nothing more to say about former.

Let us assume it uncontroversial that by the normal exercise of our sense

organs, we humans perceive the world.[1] Perceptual experience, then, is constituted by being in a certain sort of conscious mental state—the 'perceptual' sort, whatever that means—and the phenomenology of the experience is what it is like for the perceiving subject to be in that state. 'Phenomenal', then, indicates some or other relationship to phenomenology.

Phenomenology is associated with many terms in the philosophical literature– 'phenomenal character', 'qualitative character', 'qualia', 'appearance', 'phenomenal consciousness', 'subjective consciousness', 'conscious awareness', 'subjective awareness', and others. Few would accept the equivalence of all of these terms, probably most directly because many refer to states (like 'consciousness' or 'awareness'), while 'phenomenology' is an aspect of a state. Additionally, some of the terms involve special theoretical commitments, like 'qualia', or plausibly pick out something related but subtly different, like 'appearance'. In the perceptual psychology literature, for example, you will sometimes see the terms 'phenomenal' and 'phenomenology', but you also find references to the 'appearances' of perceptual stimuli, which presumably have something to do with the phenomenology associated with perceiving those stimuli but are understood to apply to stimuli and not to experiences. The relationship with awareness and consciousness is very close, but 'awareness' is probably the containing term, 'consciousness' contained within it (since, perhaps, not all awareness is conscious) and the phenomenal left as an aspect of conscious awareness. But then there are those who deny any 'phenomenal' aspect to consciousness at all (e.g., [Dennett, 1991]). To avoid confusion, I will use only the terms 'phenomenology' and 'phenomenal' and will presuppose no theoretical commitments other than what is implied in the definition

[1]When I say 'normal exercise', I mean to include all the neural processing associated with perceiving. This, surely, is part of the 'normal' exercise of these organs, even if it is a separable phenomenon in whatever sense.

above.

1.2 The Representational Contents of Perceptual Experience

These days, phenomenology (or 'phenomenal character') is often understood as one of a pair of aspects of perceptual experience, the other being its *representational content*. The representational content (or simply 'content') of an experience is the set of conditions under which it would be veridical or accurate. Perception is, by most accounts, a representational faculty—a faculty by which we represent the world around us—and so it's not unreasonable to suppose that perceptual experiences have accuracy conditions. After all, to represent the world is to portray it as *being* a certain way, and that portrayal can be more or less faithful to reality. This is, of course, a decidedly realist outlook on perception generally, and one might entertain varieties of anti-realism instead, but I will have nothing to say about the matter here.

Though I will not offer any extended argument of my own for the claim that perceptual experiences have contents, it bears mentioning some of the motivation for the claim, as this will be helpful in elucidating the notion of 'content' used here. The claim that (all) perceptual experiences have representational contents—'the content view'—has been recently defended at length by [Byrne, 2009], [Siegel, 2010], and [Schellenberg, 2011] and is often simply assumed true. The central argument for the content view is sketched above. The thought is that when you are having a perceptual experience, the world seems to you to be a certain way. The experience is not neutral about what the world is like. It 'says' that the world is *this* way, ruling out

incompatible alternatives. This way of putting the point seems equivalent to saying that experiences have accuracy conditions, since in 'saying' that the world is a certain way, perception risks error.[2]

A related point concerns hallucination. I can imagine having the same experience on two different occasions, one experience hallucinatory and one not. But if one is hallucinatory and the other is not, then the individuation of perceptual experiences cannot pertain to their *relata*—the perceiving subjects and the perceptual stimuli—as contemporary opponents of the content view claim.[3] *Ex hypothesi*, the experiences are the same, but their *relata* are not. We are therefore forced to account for the sameness of the experiences some other way, and the content view will say simply that the experiences have the same contents. Say, for example, that I have an experience as of a red globe on the table. Such an experience is possible, it is assumed, whether or not there is a red globe on the table. The globe might be another color, or the object might be a non-globe, or there might be nothing on the table at all. But whether or not there is a red globe on the table, I can have an experience as of a red globe on the table, and what makes the experience identifiable in this way, from one instance to the next, regardless of what the world is like, is that it has a certain *content*—viz., *that there is a red globe on the table*. The easy handling of illusion is thought by many to be the greatest advantage of the content view and the greatest shortcoming of 'relational' views. To be sure, these motivations for the content view, though compelling, are not decisive, and I have given no voice to the opposition. The battle between the content

[2]I use quotes around 'saying' because I am using an analogy with language to describe the way in which perception presents the world as being. I actually think the analogy with language is misleading, but it is innocuous here.

[3]More precisely, proponents of the 'relational' view of perception make this claim. See, e.g, [Campbell, 2002, Travis, 2004, Brewer, 2006]. The content view has other opponents as well—notably, sense-data theorists like [Russell, 1912] and other British empiricists reaching back to Thomas Reid, but they are not contemporary.

view and the relational view rages on, and I will not try to settle the matter here.[4]

Even among those who accept that perceptual experiences have contents, views about the nature of the contents vary widely. There are, for example, debates about whether perceptual contents can involve particulars, about whether they are individuated by conceptual modes of presentation, about whether they can involve indexicals, etc. I'll not say a great deal about these debates apart from what emerges naturally in the discussion. In my discussion of specifically phenomenal contents, I will stake out a position in the debate over the 'admissible contents' of perception, which sounds perfectly general but often focuses on what sorts of properties can figure in perceptual contents. Contemporary debates typically focus on objective, physical properties, and in my discussion, I will do the same, although I will nowhere argue that perception is limited in this way.

It suffices for now to say that I am assuming that perceptual experiences have representational contents, and that such contents are *attributional* in the sense that, at their core, they involve attributions of properties to objects. Though I won't argue for the claim generally, I will argue against one form of non-attributional content in Chapter 2.

1.3 Intentionalism

Among those who accept that experiences have contents, it is generally believed that such contents are closely connected with phenomenology. It may be thought, for example, that the distinctive feature of perception—a feature that distinguishes it

[4]It is worth noting that although I list Schellenberg as a defender of the content view, she is more usefully described as attempting to make peace between the content view and the relational view. So perhaps an armistice is coming. In any case, she does argue that experiences have contents.

from other sorts of mental representation—is phenomenology, and that therefore if an experience has representational content, it is very likely to be closely connected with its phenomenology. Thus, if we accept that experiences have content, we are likely to accept that perceptual content is bound up with perceptual phenomenology in some way.

The strongest claim about the relation between phenomenology and content is that the phenomenal character of an experience is identical with certain of its intentional content (i.e., representational content). This is sometimes called 'strong intentionalism' (or 'strong representationalism'), and despite the severity of the claim, it is not unpopular.[5] Weaker versions of the thesis are usually cast as supervenience claims, and what we might call 'minimal' intentionalism is the claim that the phenomenal character of an experience supervenes on its intentional content (or certain of its intentional contents).[6] The phenomenal character of an experience supervenes on its intentional content if and only if, *necessarily*, any two experiences alike in intentional content will be alike in phenomenal character. Or, equivalently, any adjustment to the phenomenal character of an experience necessarily demands a corresponding adjustment to its intentional content. If character and content are thought to supervene on each other, this relationship is sometimes described as one in which, necessarily, character and content *covary*.

[5]Strong intentionalism is defended, e.g., by [Tye, 1995], [Dretske, 1995], and [Lycan, 1996].
[6]See, e.g., [Byrne, 2002] for a defense of the supervenience thesis.

1.4 Phenomenal Content

Perhaps the weakest claim we'd make concerning the relation between character and content is that, necessarily, states with phenomenal character (of a certain sort, perhaps) have intentional contents. The claim is rarely defended, since nearly all parties in the debate over intentionalism accept it.[7] A slightly stronger claim than this—the claim that necessarily, any two experiences alike in phenomenal character will share intentional content—is commonly understood to mean that there are *phenomenal contents*.[8] In the language of supervenience, phenomenal contents are those intentional contents that supervene on the phenomenal characters of perceptual experiences. If we accept minimal intentionalism and the claim that there are phenomenal contents, then phenomenal content and phenomenal character supervene on each other, or covary.[9]

The claim that there are phenomenal contents is defended at length by [Siewert, 1998] and [Horgan and Tienson, 2002]. It is important to note that since supervenience is a *necessary* relation between two sets of properties, it is not accidental, on this view, that states with certain phenomenal characters have certain intentional contents. Colin McGinn is often quoted in this context as saying that intentional content is "internal to phenomenology" [McGinn, 1991, pp. 35-6]. Because of this, arguments in favor of phenomenal content cannot merely *point out* that phenomenally-similar states have similar intentional contents, because a single

[7]Even staunch opponents of intentionalism, like [Peacocke, 1983] and [Block, 1996] appear to accept it. Notice, however, that accepting this claim is tantamount to accepting the content view, since perceptual experiences are states with phenomenal character. And since the content view has opponents, the present claim does, too.

[8]See, e.g., [Tye, 1995, p. 137], [Chalmers, 2006, p. 50], [Thompson, 2006, p. 76], and [Bayne, 2009, pp. 386-7].

[9]Note that although this claim falls short of saying that the two are identical, it does not rule it out. Thus, the claim that phenomenal content and phenomenal character supervene on each other is compatible with strong representationalism.

counterexample would render the thesis false. Arguments for phenomenal content need to show that the having of a certain phenomenal character *guarantees* the having of a certain intentional content. Siewert's arguments are extensive and several, and we will not reproduce them here (for reasons I'll give shortly), but his general strategy is to argue that perceptual experiences have accuracy conditions *because of* their phenomenal characters [Siewert, 1998, p. 221]. The claim that perceptual experiences have accuracy conditions is just the content view, and we motivated the content view by saying that in perceptual experience, things seem to *be* a certain way. The claim is quite plausible, but we might ask *why* things seem a certain way. What is it about experience *in virtue of which* things seem to be a certain way?—it is, Siewert thinks, the phenomenal character of the experience. Whatever the conditions that render a perceptual experience accurate or inaccurate, those conditions depend essentially on what it is like to have the experience. Once you've imagined an experience with some particular phenomenal character, you've *thereby* imagined an experience with some particular set of accuracy conditions. Moreover, if you imagine another experience with the same phenomenal character as the first—if having the one experience is just like having the other—then it seems almost undeniable that the two experiences represent the world as being the same way. Or consider supervenience from the other direction. Imagine two experiences that represent the world as being nearly the same, but different in one detail. Imagine what it is like to have those experiences. It again seems undeniable that we are imagining experiences with different phenomenal characters. We might—if we are strong intentionalists—go as far as to say that the world's seeming to you to be a certain way *just is* its having a certain phenomenal character. If this is true, then the two are not merely contingently related.

Horgan and Tienson make similar points using the notion of 'phenomenal duplicates' [Horgan and Tienson, 2002, pp. 524-5]. They argue that any two possible phenomenal duplicates, regardless of life history or environment, will share certain intentional contents, of a type that is 'pervasive' in our mental lives.[10] We can see that this approach is quite similar to Siewert's, or at least, to our elaboration of it, except that it describes multiple perceivers instead of a single perceiver. This is important, too, because it reminds us how strong a relation supervenience really is. It is not simply that any time I have an experience with some particular phenomenal character, I will represent the world as being a certain way; rather, it is that *any* experience with the relevant phenomenal character, in this world or any possible world, at any time or place within that world, in any head, regardless of life history or environment, regardless of the underlying perceptual machinery, will represent the world as being *exactly* the same way.

I will make no attempt to establish whether or not there is phenomenal content in this very strict sense. Certainly, the considerations adduced above are not enough to establish a relation like supervenience between character and content. Fortunately for the present project, it doesn't matter whether there are phenomenal contents in this sense. For, as I will now explain, supervenience is both too strong and too weak a relation for my purposes.

[10] If we believe that phenomenal character is not fixed, even partially, by conditions 'outside the head'—if, i.e., we are internalists about phenomenal character—then Horgan and Tienson's argument is designed to establish not only that there are phenomenal contents, but that there are 'narrow' phenomenal contents. Internalism about phenomenal character is popular, though not universally accepted. See, e.g., [Byrne and Tye, 2006].

1.5 Explanation

Supervenience theses are typically put forward in reductive projects, in order to show a kind of dependence between sets of properties. And this is precisely the motivation for many intentionalists, who wish to reduce phenomenality to intentionality, as it is commonly thought possible to give a naturalistic theory of intentionality. My aspirations here are not to reduce but instead to *explain*. To be sure, true reduction yields explanation of a sort, but explanations needn't be reductive in order to be informative or important. Moreover, to establish that one set of properties supervenes on another is not enough to explain the one in terms of the other.

1.5.1 Supervenience Is Too Strong

I will not endeavor to establish either that character supervenes on content (intentionalism), or that content supervenes on character (phenomenal content). I think that the considerations typically offered in support of these theses fall short of establishing supervenience. This is not to say that I believe supervenience here to be false; rather, it is simply that I have no special *need* to show it true, given that my project does not aspire to reduction. Moreover, anyone attempting to establish supervenience confronts a mass of potential counterexamples, and intentionalists have spent much of their time responding to these, making modifications to the view, responding to strengthened forms of attack, and so on. Not only has the dialectic become tiresome, but it has shifted attention from the fact that there are rich possibilities of explanation in the basic observations that motivate the intentionalist view.

This fact has not been ignored entirely. Frank Jackson, in a paper that marks

his (rather famous) capitulation to physicalism—and to intentionalism in particular—discusses some features of perceptual representation that make it special and that help explain the phenomenal aspects of perceptual experience in terms of its contents.[11] Jackson says that the relevant representations are rich, and in fact 'inextricably' rich, meaning that the richness of perceptual representation is not separable in the way we might expect 'cognitive' representations to be. For example, when you see an object, you cannot see only its shape, or only its color. You must see its shape and color together. Because this is not true of other sorts of representations, like thoughts, perceptual representations are special. Such representations are also 'immediate' in the sense that they convey their contents without a transmission medium.[12] Jackson goes on to describe additional features, but it will suffice for our purposes to notice that what he's doing may enable us to see the special relationship between perceptual representation and its phenomenal aspect. In fact, Jackson thinks that we can *deduce* the phenomenal aspect from a suitable description of the intentional contents (a description we can give at least in principle). Once we've given such a description, Jackson argues, "we get the phenomenology for free"[Jackson, 2003, p. 270]. To be sure, Jackson has reductive aspirations. But they needn't take center stage here,

[11][Jackson, 2003, pp. 269-70]. Jackson's capitulation to physicalism is famous because he is the author of the original form of what has come to be called the 'Knowledge Argument' against physicalism, or, colloquially, the 'Mary' argument [Jackson, 1982].

[12]We might view this remark as pointing to what is called the 'transparency' or 'diaphanousness' of perceptual experience—the observation, made originally by G.E. Moore, that when we try to access the qualitative aspects of visual experience by introspection, the best we can do is to see the associated properties instantiated in the scene [Moore, 1903]. I cannot see blueness *as* a qualitative feature of my experience, I can only see that some object is blue. For obvious reasons, this intuition has been a major driver for intentionalism, since it may seem to show that all phenomenal character is intentional in nature. I have very consciously avoided appealing to 'transparency' anywhere in my own arguments, not only because the intuition itself is difficult to assess, but also because the matter of its theoretical implications is somewhat vexed. However, observations rather like the 'transparency' observation are active, I think, in any argument that tries to establish the intentional nature of phenomenal experience. For critical discussion, see [Stoljar, 2004]

because regardless of whether, for example, the phenomenal facts are deducible from the intentional facts, we've made progress at explaining the peculiar features of perceptual phenomenology in terms of perceptual content, or *vice-versa*.[13]

Now, it may be objected that the relevant sort of explanation is available only on the condition that reduction is possible. In giving the explanation, we needn't focus on reduction, but in order for the explanation to make sense, we must keep the faith. This is a serious worry, in one sense. Say, for example, that, on considering counterexamples to intentionalism, I become enamored with the suggestion that perceivers from different animal species represent the world as being the same way despite having very different sorts of experiences. Or, similarly, suppose that I become convinced of the possibility of 'spectrum shift' in human perceivers, such that the way some object looks, phenomenally, to me, is *ever so slightly shifted*, along the same phenomenal spectrum, from the way it looks to you, despite the fact that we both represent the object veridically. If these sorts of cases are possible, intentionalism fails. We might, in response, adopt a *context-relative* version of intentionalism according to which phenomenal character and content supervene on each other only within certain contextual parameters. In effect, this expands the subvenient base to include the relevant features of the context (perceiver 'type', or whatever). Now, phenomenal character does not supervene on content—it supervenes on content *plus* context. Metaphysically, there is nothing wrong with this relation as it stands. However, this revision of the thesis undermines the reductive aspirations of intentionalism. Reductive intentionalism says that phenomenal character is intentional *by its very nature*, but if phenomenal character supervenes not on content but on content-plus-context, then

[13]For a more detailed treatment of the same sort, see [Hill, 2009, pp. 90-94].

there exist one-many relationships between the phenomenal bits and the intentional bits of perceptual experience, and the phenomenal bits appear *not* to be intentional *by their very nature* and instead only by the specific role that they play in some perceptual system. And if the same phenomenal bit can play different intentional roles in different heads, it does not carry its intentionality around with it. It is, instead, a representational *medium*, or as Block says, 'mental paint'.[14]

Now, why is this problematic for a non-reductive explanation of character in terms of content? The idea is this. The sort of explanation I have in mind would involve noticing that certain features of the *explanandum* can be found to correspond with similar features in the *explanans* (and this can go both ways) and concluding that these features are such-and-so *because* these other features are such-and-so. Later stages of the process move beyond noticing shared structure to choosing amongst descriptive alternatives in one domain by adverting to corresponding parts of the other. The process of discovery of parallel structure in seemingly disparate domains is quite common in cognitive development and can be very illuminating.[15] But in the everyday sorts of cases where this process is common, it doesn't matter whether the

[14][Block, 1996]. One could respond to this worry by supposing that although different perceiver-types do indeed have different perceptual phenomenology, perhaps in virtue of having different perceptual machinery, this means only that the subvening contents are finer-grained than we thought, and that each type has its own proprietary contents. This would preserve the one-one relation between character and content. But this response is problematic. First, this approach does not appear to work against Block's 'Inverted Earth' scenario [Block, 1990]. There, the perceiver and her perceptual machinery are held fixed while the *world* is adjusted, generating one-many relations between character and content. But even if we reject the coherence of the Inverted Earth scenario, the present response sees a dramatic explosion in the number of possible perceptual contents, and it makes the sharing of perceptual contents impossible across perceiver-types (and perhaps within single perceivers across time). If the relevant contents make reference to perceiver-type (say), then the failure of content-sharing is understandable, but now our theory of perceptual content looks phenomenally inadequate. If the contents make no such reference, and the relevant contents are still 'worldly' in the desired ways, then we are forced to say that each perceiver-type, in virtue of its perceptual machinery (say), has a special representational power unique to the type, a special ability to access the world unavailable to all other types. That seems wildly implausible, and in fact doesn't even respect the original description of the 'shifted spectrum' case.

[15]See [Gentner, 1983] for one way this might work.

one domain actually stands in any sort of dependence relation with the other, because we are not trying to explain the one in terms of the other or disambiguate one by reference to the other. We are simply allowing the shared structure between them to elucidate the nature of both, and any violations of this parallelism can be shrugged off. However, as applied to a philosophical explanation of phenomenal character in terms of phenomenal content, this process might seem incoherent unless we have faith that the two domains *do* indeed stand in the relevant dependence relation. Say, for example, that by our introspection on the phenomenal character of a particular class of experiences, we come to conclude something about the intentional contents associated with those experiences. It would seem that the coherence of our conclusion depends, not on *establishing*, say, supervenience, but at least on our having faith in it. Otherwise, *it wouldn't even make sense* to say that we've learned about the one domain by inspecting the other. So again, the point here is not that the claims belonging to the explanation are unjustified (they may be, of course), it's that they imply a commitment to dependence we are unwilling to make. I can't *explain* one domain in terms of another if they simply *happen* to have parallel structure. Their parallel structure must be *guaranteed*, at least by my perhaps unjustified metaphysical assumptions about the two domains, if I am going to engage coherently in this kind of explanatory process. And to be clear, I do here wish to engage in this sort of process. Am I thereby committed to supervenience, or something similar?

I don't think so, in fact. And the reason is that even if we felt compelled to contextualize the correspondence between character and content, the resulting axes of contextual variation would all of them be idle in suitably local explanations. Moreover, even in global explanations, contextual variation is unlikely to count much,

unless we imagine that the global relation between character and content is extremely disorganized. Say, for example, that I introspect my present experiences in coming to learn something about the nature of color representation (I will attempt to do this in later chapters). The kinds of sameness and difference I am likely to find in my own experience will be explained not by features of my perceiver-type, but rather by the kinds of intentional contents I'm tokening. For a single perceiver at a particular time, the *only* axis of variation that will matter for explaining phenomenal character will be the intentional axis. It's not as if, by the failure of global supervenience, my explaining some change in my perceptual phenomenology could advert to *any* of the global axes of variation. For my own experience, only some of the axes can explain change, because, in my own experience, only some can even *see* change. And even globally, across all perceivers, the richest axis of phenomenal variation is still, *by leagues*, the content axis. Now of course, this is still a bit of faith. I can't argue for it, because I myself have only ever experienced variation along the content axis—or if I've experienced any other sort of variation, I didn't know about it. I, like all people, go about my daily life *assuming* that if I witness a change in the phenomenology of my experience, I am seeing a change in the world.[16] So there'd be no thought experiment I could construct to prompt the belief that, by and large, phenomenal character fixes content the same way in every head. But in having *this* faith, I am not committed to supervenience. I can accept that phenomenal character is mental paint. I need only trust that, within the relevant contextual parameters, the only active axis of variation is the content axis. My process of explanation will not, of course, yield reduction, but it can very well yield understanding, and importantly, it can license inferences from one domain

[16]Or at least, I am seeing either a change in the world, or some artifact of my visual system, like an afterimage. But in such cases, I don't suppose I'm entering a new representational context that could account for the change.

to the other, which is of special importance to my project, as I will be arguing from facts of phenomenology to facts about perceptual representation generally.[17] To be clear, I am here acknowledging that explanation in these domains does demand some commitment to or faith in the covariation of character and content. And this does add potential confounds to our proposed explanations. However, in the normal sorts of cases, there will be only a single active axis of variation.[18] And this means that the supervenience of character on content does not mark a kind of threshold past which we can longer give any sort of explanation whatever. Rich, fruitful explanations of character and content are available at points beyond supervenience. The further away we get from an organized relationship between the two, the weaker our explanations will be, but if we limit our investigation to human perception, as I do, it is quite plausible that we keep close to maximal organization.

1.5.2 Supervenience Is Too Weak

I've said, thus far, that the operative notion of explanation here is both stronger and weaker than supervenience. It is weaker in the ways just discussed, and stronger in the sense that it is not fundamentally a matter of establishing covariation. Rather, it is an attempt at elucidating the structure of one by adverting to the structure of the other, sometimes relying on the faith that the two stand in the relevant relationship of dependence (co-dependence, perhaps). Jackson's remarks point us in the right direction, but we should say a bit more on the matter before moving on.

Jackson's approach is instructive because it endeavors to establish something

[17] It's worth noting that Susanna Siegel does the same, without any explicit commitment to supervenience [Siegel, 2010]. I will discuss her methodology in detail later.

[18] That is, there will only be one active axis of variation on the 'explaining side'. There is always variation on the 'to be explained' side, of course.

stronger than supervenience. His strategy is not simply to identify a relationship of covariation, but rather to deduce what perceptual experience should be like, given the meaning and internal structure of the relevant contents. For example, he says that perceptual contents contain a causal element, and that this is borne out in what it is like to have perceptual experiences. We experience sounds, e.g., *as* coming from sources. This is built into the experience, he thinks. I happen to disagree with the example, but his approach is admirable, because it looks not only at whether contents and characters make all the same number of adjustments at the same times, but also at *how the contents actually speak about the world*, and whether this does justice to the corresponding phenomenology. Because, for Jackson's project and for mine, it is important that the intentional aspects of character and content are a match, we have to consider both the meanings of the relevant contents and the presentational character of the experience. In chapter 2, I will discuss the matter whether phenomenal content can be indeterminate, and it will be important, during my investigation, to consider the ways in which indeterminate contents speak about the world. The structure of my investigation there will be quite like it would be, were I testing for supervenience, since testing for contextualized supervenience is just testing for supervenience within certain parameters, but unlike a test for mere supervenience, my investigation will enforce an additional constraint concerning the explanatory adequacy of the proposed relationship. Just what exactly explanatory adequacy amounts to, it's tough to say, but I will leave judgments of adequacy to the consideration of specific cases.

1.6 Perception and Cognition

In the previous section, I defended non-reductive explanation generally, but I haven't yet said why it'd be useful to give a non-reductive explanation of *perceptual experience* in particular. And so in this final section of the chapter, I'd like to say a bit about the deep motivations for the project.

Thus far, we've simply taken for granted that 'perceptual' representation can be distinguished from other sorts of representation, and specifically 'cognitive' representation. In one sense, the two are very close together, since they are both kinds of specifically *mental* representation, and probably the *only* kinds of mental representation. But why are there two?—and I don't mean what *caused* there to be two, from an evolutionary perspective. The question, rather, is why look at what we know about the brain and *declare* that they are two and not one?

We should be careful here to distinguish our question from debates over so-called 'cognitive penetration'.[19] This is not to say, of course, that the question under consideration in those debates is not worth asking. It is a perfectly good question about the computational dynamics of mentation. But it is *not* the right question to ask if you wonder whether perception and cognition should count as distinct representational systems. It is not simply the same question with richer empirical detail, for example. And this should be clear enough, since, far from of undermining the distinction between perception and cognition, the debate over cognitive penetration sustains and galvanizes that distinction—the very statement of the question of penetration presupposes it, in fact.

And thankfully, *our* question has a much easier answer—perception and cogni-

[19]For a clear statement of the thesis that visual perception is impenetrable to cognition, see [Pylyshyn, 1999]

tion are distinct representational systems, because they can disagree. In the Mueller-Lyer line illusion (Figure 1.1), for example, the lines look to be of different lengths, whether or not the perceiver believes them to be of different lengths. Learning that the lines are the same length does not break the illusion. And if there are two answers, then there are two systems.

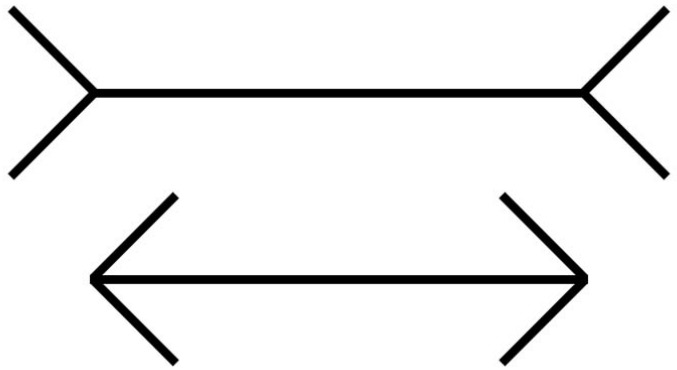

Figure 1.1: The Mueller-Lyer line illusion.

But what is perhaps more interesting than our having two different answers is that the answers are given in very different ways. There are two entirely different representational idioms here. Descartes' example of the wax in the *Second Meditation* is an excellent demonstration of this point. Descartes observes that no matter how many reformations of the wax you witness, you cannot *perceive* that the wax is changeable. For, changeability implies more than what can be seen or felt. Through the senses, we are limited to seeing the wax move through some finite subset of its possible formations. But the notion of changeability vouchsafes an infinity of perceptual experiences, and nobody can have an infinity of perceptual experiences. Despite this, we can represent things as being changeable, so the notion must be extra-perceptual. Because there is no experience or set of experiences that can convey

the content, perception cannot token the content. There are two things to notice here. The first is that perception is *limited* in ways that cognition is not. There are some contents simply off-limits to perception. Arguably, this is not true of cognition, at least if we focus only on the extensional aspects of content. Second, perceptual representation is limited by what is possible *phenomenally*. This claim is stronger than the claim that there is phenomenal content. It is the claim that *all* perceptual content is phenomenal content. If this is true, then it should be possible, by examination of perceptual experience, to identify a set of constraints on perceptual representation. I intend to do precisely that, in view of some particularly difficult puzzles of perceptual experience. The overarching constraint, as we shall see, is that because phenomenal experience is *presentational* and not *descriptive*, perception is limited to representing what can be presented. Perception shows; it cannot tell.

Chapter 2

Indeterminacy

As discussed in the previous chapter, phenomenal contents are those perceptual contents that supervene on the phenomenology of perceptual experiences, such that, necessarily, any two experiences alike in overall phenomenology will be alike in overall phenomenal contents. It is my concern in this chapter to argue about such contents that they cannot be indeterminate. I will do this by endeavoring to show that indeterminate contents cannot explain the phenomenology of perceptual experiences. But we'd like to explain the phenomenology of perceptual experiences in terms of their contents. We'd like to say, e.g., that two experiences are phenomenally alike *because* they represent the world as being the same way. Accordingly, I will argue, the relevant contents cannot be indeterminate.[1]

Before I proceed, a word about the language I will use here. As I discussed in the previous section, the notion of explanation at work here is both a weaker and a stronger relation than supervenience. Explanation in the sense I mean demands

[1]Ned Block has a view according to which the relevant contents are indeed indeterminate; he concludes that such contents therefore cannot determine perceptual phenomenology. See [Block, 2010]. Thus, in one sense, we are in agreement. However, where Block is willing to give up explaining experience in terms of content, I am not.

only *contextualized* and not global covariation between phenomenal character and content. I suspect that the relevant contextual parameters are types of perceivers (and so, usually, heads) and times. Because of this, I think it appropriate to refer to the fixing of these parameters as a kind of 'localization'. Within local contexts, I will suppose that character and content necessarily covary, and thus 'supervene'. I will therefore refer to the relevant relation as 'local supervenience', or l-supervenience. A set of properties A l-supervenes on a set of properties B if and only if, *locally*, any A difference necessitates a B difference, where 'locality' is understood as the fixing of the relevant contextual parameters, whatever they are.

Additionally, because I think l-supervenience does not suffice for explanation (since it is weaker than supervenience and supervenience doesn't suffice), I will assess the relevant cases for more than just l-supervenience; I will, additionally, assess them for explanatory adequacy. I will not attempt to give anything like a definition of explanatory adequacy, but I can say that the focus of my assessment will be the ways in which phenomenal characters and contents are directed at the world—i.e., their intentionality. As such, it will matter a great deal, e.g., how any proposed intentional content of experience speaks about the world and how any experience under consideration presents it. As I've said, I think that the ways in which perceptual contents can speak about the world is limited by the ways in which perceptual experience can present it, but I mean to draw out that consequence by my discussion, rather than simply to stipulate it from the outset. Think of it as a working assumption to be vindicated by my investigation.

Finally, as I said in the previous chapter, I will nowhere claim that there are phenomenal contents in the strict sense, since I think that the reasons offered in favor

of their existence do not suffice for the supervenience of content on character. However, I *do* think that they typically suffice (if true) for *l*-supervenience. Accordingly, I will use the term 'phenomenal content' to mean that intentional content that *l*-supervenes on the phenomenal character of perceptual experience.

2.1 Clarification of the Thesis

In saying that phenomenal contents cannot be indeterminate, I do not mean that phenomenal contents cannot be determinable with respect to some more determinate content elsewhere in the system. Because properties stand in various determinacy relations with one another, representational contents that involve properties can stand in analogous relations to each other. Red is determinable with respect to scarlet; analogously, we can say that the content *red* is determinable with respect to the content *scarlet*. However, that *red* stands in this relation to *scarlet* says nothing about whether some particular *red* token is, we might say, 'locally' determinate—i.e., determinate with respect to its casting in the determinacy hierarchy.[2] A system can token determinate *red* even though it is capable of tokening *scarlet*. I doubt there is any way to identify a maximally determinate property, but in any case I am not claiming that phenomenal contents involve only such properties.

I mean when I say that phenomenal contents cannot be indeterminate that they are subject to something like the law of excluded middle. A property P and object o can figure in the phenomenal contents of an experience only in the form $P(o)$. If o is not represented as having P, then o and P will not figure in the phenomenal contents of the experience. There is no counterpart in phenomenal content for Bayesian 'degrees

[2]Positions in the 'determinacy hierarchy' can be understood as levels of precision.

of belief', for example. Any model of representational content according to which o can have P with less than 100% confidence is inappropriate as a model of phenomenal content.[3] Or so I will argue.

Now, it is possible in one sense for the perceptual system to token a content describable, albeit misleadingly, as a range of values. For, a range of values can be extensionally equivalent to a single property. There are, for example, many shades of red, and we can describe redness, then, as a range of red shades. Extensionally, then, visual contents can be equivalent to ranges of values. And this might make such contents seem 'indeterminate'. However, I think that ranged constructions threaten my thesis only if they are understood *intensionally*—only if, i.e., the ranged contents are presented *as* ranges in experience, and this because the presentation of a range *qua* range is an admission of ignorance about the thing represented, and so counts as an attributional failure at its casting level. To say that the visual system represents the color of an object as a range of values is to say that the visual system has chosen to represent an object at position p in the determinacy hierarchy, but it is not sure which value at p is true of the object. To say, for example, of a man that his height is somewhere between 60 and 70 inches is to say that you don't know his height in inches. It may be said that you know his height in decade-inches, but you do not know

[3]James Stazicker calls the inadmissible contents 'vague' and the admissible ones 'indeterminate' [Stazicker, 2011, p .170]. I mention this because in my terminology, the inadmissible contents–the ones that cannot explain perceptual experience–are 'indeterminate', and the admissible ones are merely 'determinable'. In truth, the matter is somewhat more complicated than this. Stazicker argues in favor of visual indeterminacy, but his proposal is compatible with what I say here, because my arguments are about the *intensional* aspect of the representation, and his are only about its extensional aspect [Stazicker, 2011, p. 173, n. 13]. See also [Stazicker, 2014, p. 6, 21]. I will concede below that there are visual representational contents extensionally equivalent to ranges of values and that are therefore in Stazicker's sense 'indeterminate'. What I deny (*inter alia*), and what Stazicker likewise seems to deny, is that perceptual experiences can fail to have determinate accuracy conditions. He calls these cases 'vague', but I think that language is too narrow for my purposes, because it calls to mind, primarily, borderline cases of class membership, and my focus is on perceptual comparisons.

his height in inches. The visual system, then, to put the point in the terms just used, can say of an object that it is *sixties decade-inches* tall, but it cannot say that it is *60-70 inches* tall, because the latter construction is a p-level representation indicating indeterminacy at p. Whereas, the former construction is a p-1 level construction indicating determinacy at p-1. And the fact that it implies indeterminacy at p is unimportant, because as I've said, I am not denying that perceptual experiences fail to pronounce (or fail to pronounce determinately) at some positions in the determinacy hierarchy. They do so fail, without a doubt. But again, that implies nothing about their local determinacy. Later on, I will argue that the *sixties decade-inches* construction better reflects our phenomenal judgments about objects.

There is the further question whether phenomenal contents are locally *maximally* determinate. Say that I am correct that no phenomenal contents are indeterminate. It might still be possible for the visual system to 'hold back'—to represent, though determinately, at p at location l, despite being capable of representing at $p+1$ at l. I don't know whether this is possible, but it's not important here. Maybe the visual system will do something like this in budgeting its representational resources, for example. The important point here is that whatever p, the representations at p will be determinate.

My view presupposes, then, that representational contents are assignable to positions in the determinacy hierarchy. And we might ask what decides some content's position in the hierarchy. Causally, of course, it is determined by facts about the sensory system in question—e.g., its acuity. But analytically, a representation 'a is R' belongs to the position in the hierarchy occupied by Rs. That is, if you have a representational content involving R, and R is at p, then the representation is

assigned to position p. This may seem strange at first, since in representing R, your representation carries information about properties G and S related, nomically or analytically, to R. And G and S might very well reside at positions higher or lower in the hierarchy. Thus, we might wonder which among R, G, and S decides the position p of the representation. It is part of the peculiar nature of intentional representation, however, that intentional content is unlike informational content in some ways. One among those ways is that intentional contents do not convey all of the information analytically or nomically 'nested' within [Dretske, 1981, p. 173]. Although such information might be carried elsewhere in the system, it is not represented in the intentional contents of that subsystem.

2.2 Phenomenal Content Cannot be Indeterminate

In this section, I will begin with the assumption that we'd like to explain the phenomenology of perceptual experiences in terms of their (phenomenal) contents. And this desideratum seems reasonable on the face of it. If I consider two perceptual experiences as of a stucco wall [Fig. 2.1], phenomenally alike except in the fine details of texture, I will assume straightaway that what explains this difference is a difference in the way my experience represents the stucco as being—in this case, likely a difference in the way the stucco was applied to the wall to create its particular topography.

This is explanation in the substantive sense—we do not say simply that the two experiences differ phenomenally and in terms of their contents in all the relevant cases. Rather, we say that the two experiences differ phenomenally *because* they differ in their contents. In what follows, I will show that indeterminate contents cannot substantively explain the phenomenology of perceptual experience. I will

Figure 2.1: Stucco walls. We experience these differently because we represent them as *being* different.

argue that both the l-supervenience of phenomenology on phenomenal content and the l-supervenience of phenomenal content on phenomenology (which l-supervenience is part of its definition) are necessary conditions for explanation. I will then argue that although it may be possible for phenomenology to l-supervene on indeterminate contents, l-supervenience fails in the other direction. And since both directions are necessary for explanation, explanation fails. Moreover, I will argue that indeterminate contents cannot substantively explain perceptual phenomenology because they leave open possibilities that are closed off by the phenomenology. Thus, even if the argument from l-supervenience fails, the argument from explanatory adequacy shows that phenomenal content cannot be indeterminate.

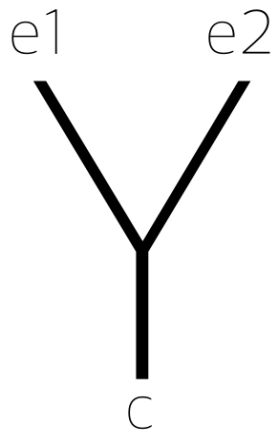

Figure 2.2: Fork. In this structure, c does not determine either e.

2.2.1 The *L*-Supervenience of Phenomenology on Content is a Condition on Explanation

In Figure 2.2, c is a content (or set of contents) and the es are experiences. If experiences did not l-supervene on contents, then the above branching structure would be possible. You could have any number of experiences associated with the same content. And it is obvious that this rules out the content's explaining the experience. For, one question you'd like answered, were you to explain the one in terms of the other, is this: 'Why am I having experience e_1 instead of experience e_2?' And if the above branching structure is possible, then simply pointing to c cannot settle the matter. In this structure, c can lead to e_1 or e_2. Thus, we cannot say that I am having e_1 because I am tokening content c. At least, that wouldn't be enough. The tokening of c may very well constrict the space of options for experience, but it doesn't decide between e_1 and e_2.

Another way to understand the need for this direction of *l*-supervenience is to consider the making of highly precise perceptual comparisons, something we do regularly—so regularly, in fact, that any skeptic about the systematic accuracy of such comparisons bears the burden of proof here. Consider viewing two objects, very much alike in appearance. In fact, imagine that the two objects are very nearly the same color, such that the difference in color between them is the smallest difference detectable by our visual system. In such a case, we trust perception to deliver at least the verdict that the two items are differently colored. And we decide that the items are differently colored, say, only on the strength of our experiential evidence. The experience of looking at one is ever so slightly different than the experience of looking at the other (and the difference is in the color appearance of the surfaces). If the experiences in these cases did not *l*-supervene on sets of representational contents, then there'd be no guarantee that the visual system is representing the two items as differently colored. For experiences to *l*-supervene on contents is for any difference between experiences to be accompanied by a difference in contents. *Ex hypothesi*, the experiences are different, and if the contents are not different, then *l*-supervenience fails. Thus, the *l*-supervenience of experiences on contents is required by our trusting very ordinary perceptual comparison strategies. That, I think, is clear enough.

2.2.2 Phenomenology Cannot Explain Indeterminate Contents

Perceptual experiences cannot explain indeterminate contents. Consider a case where two experiences differ phenomenally with respect to the vertical retinal size of a

pair of objects.[4] If height properties are attributed determinately to the objects, then we should expect that the contents differ between the two experiences—the objects do not appear to be the same height, after all. If, however, height properties are attributed indeterminately to the objects in the form of ranges from 60 to 70 inches in one case and 62 to 72 inches in the other, then for all the experience is telling you, the objects could very well be the same height! In this way, the contents of the experience leave open an alternative that is closed off by its phenomenology.

It is worth pointing out that *l*-supervenience requires for any phenomenal difference only a difference in content, and on the assumption that the relevant contents are structured as ranges, difference in content is just difference in range, despite the fact that by identifying a range, you imply that the object can occupy any one of the positions within its boundaries. In any case, it seems that at least in some structural sense, phenomenology can *l*-supervene on indeterminate content, and so indeterminate contents seem to meet this minimal criterion for explanation.[5] But the *sixties decade-inches* construction better reflects phenomenal experience than does its extensionally equivalent, ranged counterpart, because in any case where object-phenomenologies fail to match, the experience is apparently ruling out sameness between the two objects, at least on the variables of interest—whatever properties make up the object's visible aspect. Ranged constructions build into perceptual experiences a kind of scientific wisdom they seem not to have. On the range view, the perceptual system reaches beyond its own representational capabilities and says,

[4] Assume that the distance from the viewer is fixed.

[5] I am hesitant to say with authority whether *l*-supervenience succeeds here or not. The difficulty here is that there is little precedent for assessing supervenience relations between anything but sets of properties. It's not clear to me which things even count as properties in the present investigation. In any case, I am conceding to the indeterminacy camp that it is not by failure of *l*-supervenience that their contents fail here.

cautiously, 'I have estimated the values of these two objects on the variable of interest, and although the estimates differ, because they mark out only the boundaries within which the actual values can fall, it may very well be the case that the objects do not differ on the variable of interest'. So, despite the system's registering a verdict at position *p*, it looks beyond to position *p+1*, and admits its inability to distinguish between stimuli at that level.

Similar points can be made about contents structured as confidences or degrees of belief. That phenomenally distinct experiences might attribute to a pair of objects two distinct 'confidence' contents—say, *scarlet with 80% confidence* and *scarlet with 85% confidence*—does not settle the matter whether the objects are the same color, a matter that does indeed seem settled by the phenomenology. That they might be the same color is not ruled out by the contents, despite the fact that, in some sense, the contents are different. And I think that our feeling that indeterminate content in these cases is inadequate to the phenomenology is just the intuition that our visual experience, for example, tells us about colors—it does not tell us about contents. That the contents are 'different' in whatever sense seems irrelevant; what we want to know, and what every last person possessed of this pair of experiences is disposed to believe, is that the objects are differently-colored.

2.2.3 Indeterminate Contents Cannot Explain Phenomenology

In the next two sections, I will reverse the order from the previous two. There, I argued first that *l*-supervenience is a condition on explanation, and then I argued that phenomenology cannot explain indeterminate contents. Here, I will argue first

that indeterminate contents cannot explain phenomenology, and then I will argue that content must *l*-supervene on phenomenology if it is to explain it. That this direction of *l*-supervenience is a condition on explanation is more controversial and so will require more argument.

I finished the previous section with an argument from comparative perceptual judgments, and it is worth pointing out that an analogous argument can be made for this direction of explanation. If you have two experiences phenomenally alike with respect to the vertical retinal size of a pair of objects, then you should expect that the experiences represent the two objects as being the same height.[6] And this is true regardless of which height properties are attributed to the objects. Maybe they are highly determinate height properties, maybe they aren't; what is important is that the experience *attributes them* determinately. Now, were the experiences to attribute the height properties indeterminately, in the form of a range—e.g., *60-70 inches*—then we must say of the experiences that they communicate the possibility that the two objects are not the same height. Experiences aside, if you *tell* me, about two objects, that their heights fall within the boundaries of a range from 60 to 70 inches, then it would be unreasonable for me to assume, given what you've said, that the two objects are the same height. In fact, without knowing anything else about the two objects, it is far more likely that they are not the same height than that they are. But we should expect that two experiences phenomenally alike with respect to the vertical retinal size of a pair of objects rules out the possibility that the two objects differ in height.

But there is another argument to be made here: indeterminate contents *cannot even* l-*supervene on experiences* because there are pairs of indeterminate contents

[6] Again, assuming we've controlled for viewing distance.

that cannot plausibly demand different experiences. And if contents *l*-supervene on experiences, then for every pair of distinct contents, there must be an associated pair of distinct experiences.

Suppose we visually represent an object as having some color indeterminately. Say, for example, that the representational content of the experience is *scarlet 30%*. Now consider that the reason we fail to assent to scarlet with 100% certainty is that there are always other candidate hypotheses.[7] To say of some object that it is *scarlet 30%* is to say more than that. It is to imply that there is a field of competing hypotheses the system has yet to choose among. In fact, we needn't even claim that the system is itself entertaining any alternate hypotheses. We need only point out that the system's tokening *scarlet 30%* on the one hand and *burgundy 40%* on the other could very well result in the same experience in either case. Consider the two contents *scarlet 30%* and *burgundy 40%*. The *l*-supervenience of these contents on experiences requires that the experiences of the associated objects be different. But reflect on your experiences of objects for a moment. Say, for example, that you can modulate one visual aspect of the experience with a knob. Say that you have a knob for the hue aspect of the experience. Imagine turning the knob through its positions. Presumably, the confidence score for scarlet will vary continuously as the knob moves through its positions. But that is also true of the confidence score for burgundy. So, for each 'experience setting', there are associated confidence scores for both scarlet and burgundy. Thus, even if the tokening of *scarlet 30%* does not involve the simultaneous tokening of *burgundy 40%*, there will be pairs of object-contents, one

[7]For the general case of *o*'s having *P*, the minimal number of alternate hypotheses is one: *o*'s not having *P*. But color properties are such that if an object fails to have one, it must have another. Many other properties are like this as well. These are often called 'substitutive' properties and are contrasted with 'binary' properties.

scarlet 30% and one *burgundy 40%* (or whatever), each associated with exactly the same perceptual experience. Thus, *l*-supervenience fails. I conclude that if contents are indeterminate in the sense I am interested in here (i.e., locally indeterminate), they cannot *l*-supervene on perceptual experiences.

Just to hammer home the point, consider what would be the experiential difference between representing one object as *scarlet 30%* and another as *scarlet 40%*. Would the hue change? The saturation? Say that the hue changes. Well, now it looks like we ought to be able use a new hue term (I'm assuming scarlet is a hue term) and assign a new confidence score to it. But now we have the same problem as before—we have two contents associated with the same experience. But maybe confidence is about saturation. What if the saturation changes? Now, we ought to be able to talk about a new shade—pale scarlet, perhaps—and assign a new confidence score. So we have *scarlet 30%* (say) and *pale scarlet 40%*. And it's no use to say that confidence scores for scarlet simply reflect the pale and deep shades of scarlet—imagine, e.g., that deep scarlet is really just scarlet with high confidence—because it seems perfectly meaningful to assign a confidence score to pale scarlet. And so again, we have the problem that a single experience is associated with a host of indeterminate contents (an infinite host, perhaps), and this is just to say that the *l*-supervenience of contents on experiences fails.

The problem, more generally, is that we cannot identify in our experience enough degrees of freedom to match those degrees of freedom added to the representational calculus by allowing confidences to figure in the contents of perceptual experiences. If contents can involve all of the properties we take ourselves to represent in perception, plus confidence scores for every such property, we now seem to have a mismatch

between the magnitude of all possible phenomenal changes and the magnitude of all possible content changes. And notice that this argument is distinct from the argument above according to which phenomenal differences seem to be differences in, e.g., surface color rather than differences in the boundaries of a range, or differences in confidence. The present argument does not depend on our intuitive sense that this or that phenomenal change signals this or that representational change.

I anticipate an objection here. I began by pointing out that for all we know, *scarlet 40%* and *burgundy 30%* will be associated with precisely the same experience. Choose an experience, modulate some phenomenal aspect of it, and confidence scores for both scarlet and burgundy should change. Thus, for any scarlet confidence score, there ought to be a corresponding burgundy confidence score. And in fact, there ought to be a confidence score for every possible color. But then it seems possible to associate each of these distinct contents with the same experience, and *l*-supervenience fails rather dramatically! An opponent might object, however, that 'confidence' contents are *sets* of confidence scores, and so, sure, *scarlet 40%* and *burgundy 30%* are associated with the same experiences—this means only that the associated experience has as its content {*scarlet 40%*, *burgundy 30%*}. Notice, however, that in order for this objection to work, contents must assign scores to every possible color, because otherwise, I can run the same argument on the sets that I ran on scarlet and burgundy. Imagine, for example, the limit set of colors (all of them), complete with confidence scores for each. Now imagine a set just like the limit set, minus a single color-confidence pair. It will be true of the two sets that for a given experience, each one of their constituent colors will be paired with a confidence value. And even if all (but one) of these pairs are shared between the two sets, that the second set is missing a single color means it

is a distinct content. L-supervenience fails. Only if each content set has a value for every possible color—or at least, every possible color distinguishable by the visual system—can the theory be sure to satisfy the l-supervenience requirement.

How would this theory look? Well, we know that the possibilities for human color phenomenology are exhausted by some three-dimensional color solid. Or what is perhaps simpler, we know that the human visual system (or the union of all human visual systems) can make some finite number of color discriminations.[8] We can arrange these discrimination points along a single axis (in whatever order, so in a strict sense it's not really an 'axis') and then arrange confidence scores on an intersecting axis. Though we define the x-axis in terms of possible experiences, it now represents a content dimension. Each content is a curve in this space. And if color contents l-supervene on color experiences, then every possible color experience is associated with its own distinct curve in this 2D space. The problem, then, is that we defined the x-axis in terms of the *complete* set of possible color experiences, and now we have a second axis that opens a new space of possible color experiences. Impossible.

But perhaps this is too quick. One might object that I've underdescribed the space of possible color experiences. An advocate for indeterminate content might simply say that the color solid is only a proper subset of the entire phenomenal color space. And if this is true, then I've understated the phenomenal degrees of freedom. This may be. In fact, I'm quite sympathetic to the claim that the color solid underdescribes color experience, a matter I will discuss in detail at 2.4.3. But I think, nonetheless, that there is no room in experience for non-attributional contents, a category which includes confidence contents. This is a direct consequence, I think, of

[8] '10 million' is an oft-cited figure here—doesn't matter for our purposes.

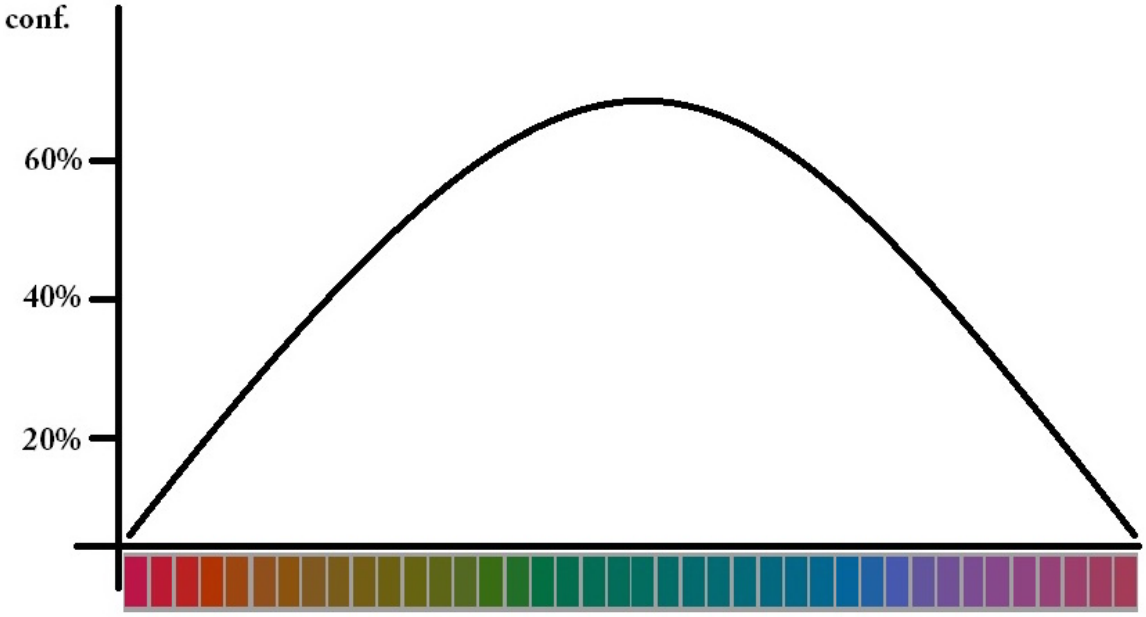

Figure 2.3: A 'color curve', which represents a single color content on the hypothetical 'confidence' theory of color.

the *presentational* nature of perceptual experience. The content 'P with x% confidence' is not *showable*. It is *describable* and *intelligible*, to be sure, but it is not showable, because it doesn't actually say *anything* about the world.[9] This sort of content cannot be an intensional match with any phenomenal character, then, because phenomenal character is always attributional. It may not always be obvious exactly how it is speaking about the world—and this because we are used to describing and thinking about the world, and not just perceiving it, and we find sometimes that perception fails to answer some of our questions about it. But that is certainly not enough to show that perception can present indeterminacy.

Consider the phenomenon of visual occlusion, for example. On the far side of the hills, there could be anything. Your visual experience of the hills doesn't settle

[9]You might object that it says something probabilistic about the world. I don't actually think this is true; it appears to say something about the perceiver, if it says anything, but in any case, probability distributions are not showable either.

the matter. But that certainly doesn't mean that, for any object possibly on the far side of the hills, your experience assigns some degree of confidence to its being there. Crucially, I'm not arguing that it *doesn't*; I'm saying only that *the mere fact* that perception can't settle a matter decisively says only that the matter does not figure determinately in the content of the experience. It says nothing else.[10] We are left, then, to decide between non-representation and indeterminate representation. 'Extensionally', we might say, there is no difference between remaining silent and saying that you don't know. In both cases, you've communicated zero information about the world. But since phenomenal experiences are themselves always determinately comparable—we can always say whether the experiences differ or not—the contents that explain those experiences must speak of the world in ways that are comparable, and here they do not.[11]

I can see a subtly different sort of objection here, however. I've shifted between talking about contents structured as ranges and contents structured as confidence scores or degrees of belief, because I think they fail in much the same way. But when we consider what I've conceded regarding ranged contents, the degrees of freedom argument may be in trouble. I've said that I think the extensional variant of the range view does not threaten my position, because I'm happy to admit that perception can

[10]It's worth pointing out that I do not run afoul of arguments to the conclusion that we can have a phenomenal experience of absence [Farennikova, 2013]. I'm not saying we can't; I'm saying only that we do not have an experience of absence for every object that might be absent. I hope that this is uncontroversial.

[11]I should say that it may look like my argument from degrees of freedom has sort of sneakily transformed into a point about explanatory adequacy. I'm willing to concede the point if need be, since dialectically, it doesn't really matter. But the way to assimilate the present point to the degrees of freedom point is simply that, once we've defined the entire set of distinguishable phenomenal experiences, there will correspond to each of them some content, distinguishable from the next with respect to its speaking about the world, and thus, to add confidence contents to this picture is again to run out of room for phenomenal variation. But as I've said, I can see how one might think all the heavy lifting in the argument is done by the criterion of explanatory adequacy.

present properties at *p* that are in fact determinable at *p+1* despite never presenting them as such. But in order to avoid running afoul of my degrees of freedom argument about color, it might seem like the visual system will need always to represent at a single level of the determinacy hierarchy. Say, for example, that the color solid is defined by all color experiences at *p*. This is a reasonable assumption. But if even a single *p+1* experience is possible, we now need another degree of freedom in the phenomenology to account for it. This is a serious problem. How, phenomenologically, is an experience of *red* supposed to differ from an experience of *scarlet*? Metaphysically, the one contains the other, so they are not competitors, but we tend to think of color *experiences* as all of them mutually excluding each other. How, then, can perceptual experience sometimes present *scarlet* and sometimes *red*?

One way to respond to this worry is simply to say that the visual system does indeed stick with a single level of the determinacy hierarchy. And I think that at least for center-field stimuli, this is true. But I am sympathetic to the suggestion that things are different in the periphery. If they are, then I must admit that the color solid account of color experience is incomplete, and that there is at least one additional dimension of color experience, one that could perhaps make room for indeterminacy in the contents of that experience. My considered view, in fact, is that colors are indeed representable at different levels of the determinacy hierarchy, and there are associated phenomenologies at each level. However, I nonetheless maintain that the attribution of color, whatever its determinacy level, is determinate. Even for peripheral stimuli, our experiences do not leave it open whether objects have this-or-that color at the local determinacy level. I will consider the case of peripheral stimuli in more detail in section 2.4.3, where I will discuss a well-developed proposal involving indeterminate

contents (specifically, confidence contents).

2.2.4 The *L*-Supervenience of Content on Phenomenology is a Condition on Explanation

We've established that indeterminate contents cannot l-supervene on perceptual phenomenology. So, why is this direction of l-supervenience a condition on explanation? Consider that if it is possible that in virtue of the experience e you're having, a has either R or Q—that is, if the relevant contents fail to l-supervene on experiences—your perceptual experiences cannot decide between R and Q. It will be true of R and Q that they are phenomenally indistinguishable. And moreover, there will be some other property or cluster of properties that both objects have in virtue of being phenomenally indistinguishable. Rather than saying, of your perceptual experiences as of Rs and Qs, that they are the way that they are because the objects in question are Rs and Qs, you ought to say that the experiences are the way that they are because the R and the Q have property P in virtue of which they are phenomenally indistinguishable. Maybe P is a shape property, or a color property, or some combination of these, but whatever it is, it better captures what your perceptual system is hooking onto than do either R or Q. You cannot say, for example, that your perceptual experience is telling you that something is an R. Why? Because it might be a Q, *even in optimal perceptual conditions*. I add the italicized clause because I anticipate the following objection: to say that e cannot represent that a is R because a might be a Q—i.e., because a might not be an R—is to eliminate the possibility of error. And that's a problem. Fair enough. But I am here considering a case in which the perceptual conditions are optimal and perception nonetheless cannot distinguish between two

objects because their perceptible profiles are identical. And this, we can all agree, is possible even if the two objects belong to different classes. If Q and R are such classes, then perception cannot tell you that something is an R.[12] *A fortiori*, you cannot say that your perceptual experience is the way that it is because it is telling you that a is R.

Now, you might say that your perceptual experience is the way that it is because it is telling you that a is *either* R *or* Q, but if this is true, it can only be because Rs and Qs are phenomenally indistinguishable. And then you can identify whatever property or properties make Rs and Qs phenomenally indistinguishable. Whatever the properties, they will be shared by both objects in question, because it is precisely the having of those properties that makes the two objects phenomenally indistinguishable. Thus, the properties in question cannot be R and Q, because they are not shared.

Let's consider an example. Say that we explain visual experience e by noting that e says that a is a duck. Surely this places some constraints on the experience. But notice that we could have the very same experience e if we were looking at a suitably crafted decoy duck. Ducks and decoy ducks are visually indistinguishable.[13] And this makes explaining e in terms of ducks seem odd: there is extra information that comes with something's being a duck that vision cannot possibly care about.

[12]Or, if you like, vision cannot capture classes R and Q phenomenally. I can see an objection here: that vision cannot distinguish between Q and R does not imply that it will never token R. It may have no business tokening R, since there are Qs out there and they look just like Rs, but all this means is that vision will *always* token R, regardless of whether the object is an R or a Q. This is an interesting problem I will discuss later.

[13]This is true regardless of whether one is a duck naif or a duck expert. If one is a naif, neither experience can possibly involve a tokening *duck*; if one is an expert, then assuming the decoy is suitably constructed, *both* experiences will token *duck*. Thus, they are visually indistinguishable, at least for a single perceiver at a time. The question whether duckhood can figure in the contents of visual experience is controversial, and I will consider it later.

Vision doesn't pick out ducks; it picks out duck-looking things, and the properties we use to flesh out *duck-looking* ought to be the ones that explain visual experience. Explaining visual experience using properties like duckhood is problematic in two directions: one the one hand, something's being a duck places far too few constraints on visual experience to explain it completely, since not all ducks look the same; and on the other hand, its being a duck involves much that is irrelevant to visual experience, like, e.g., being born of ducks. In this way, duckhood is an explanatory mismatch. There are other properties, however, that on the one hand do all of the explaining we need them to do, and on the other hand, do nothing more than that. These are the properties I argue explain visual experience (and *mutatis mutandis* for other modes of perception). One thing to notice about such properties is that they must be visually distinguishable, on pain of susceptibility to the above argument concerning R and Q. Because I think the meaning of 'phenomenally distinguishable' is likely to be misunderstood (and perhaps thought analytically true of all phenomenal content), I will call 'phenomenally-bounded' any content that satisfies the constraint discussed here.

If this reasoning is correct, then on the assumption that we'd like to explain perceptual experiences in terms of certain representational contents, those contents cannot be indeterminate. For, on the assumption that they are indeterminate, they cannot l-supervene on perceptual experiences. And if they do not l-supervene on perceptual experiences, they cannot explain them.

2.3 Siegel and the Method of Phenomenal Difference

Two claims have emerged from the discussion thus far. The first is that phenomenal contents—those contents in terms of which we'd like to explain perceptual phenomenology—cannot be indeterminate. The second is that phenomenal contents must be comprised of properties that are phenomenally bounded. Both claims issue from the desire to explain experiences in terms of (certain of) their contents. The latter claim—that the constituent properties must be phenomenally bounded—appears to rule out properties like duckhood. I turn to this matter now.

Susanna Siegel argues that natural kind properties like duckhood can indeed figure in the phenomenal contents of (visual) experiences [Siegel, 2010]. Since I am here concerned primarily to maintain that the contents that explain experiences must involve experientially (phenomenally) bounded properties, it is worth pointing out that my disagreement with Siegel is relatively minor in the sense that she needn't reject this claim. She does reject one plausible application of this claim, but her belief that, e.g., the property duckhood figures in the phenomenal contents of experiences—for her, visual experiences—is compatible with the claim that duckhood makes a difference to experience. In fact, her argument depends on the assumption that it makes such a difference.

2.3.1 The Method

At the core of Siegel's argument is what she calls the 'method of phenomenal difference'. The method begins with the presupposition that a pair of experiences

differ phenomenally. She applies the method to a number of cases, but we are here concerned only with the case of natural kinds. The relevant pairs of experiences always differ in one important respect: the subject of the experience is in exactly one of the cases possessed of a certain recognitional capacity. Call this subject the 'expert' and the other the 'naif'. In both cases, the subject of the experience is viewing an object that belongs to a natural kind, and the expert's recognitional capacity is for members of that kind. On the assumption that the two experiences differ, we can ask what explains the difference. Siegel considers various candidate *explanantia* but concludes that what explains the experiential difference is a certain difference in contents. Predictably, the expert experience involves the natural kind and the naif experience does not. In what follows, I would like only to resist the claim that the relevant contents are in fact natural kinds in the scientific sense. I will concede that experts can possess recognitional capacities the exercise of which makes a difference to phenomenology. And I will show that this admission is compatible with saying that the contents that explain experiences must be phenomenally distinguishable.

2.3.2 Why Not Ducks?

At 2.2.5, I argued that duckhood cannot be among the properties that figure in the phenomenal content of experiences. And since Siegel restricts her discussion to visual experience, I will here do the same. In short, I argue that duckhood of itself makes no difference to phenomenology. Therefore, duckhood is not a phenomenally bounded property, and it therefore cannot figure in the contents that explain visual experiences. It is important to remind ourselves that although Siegel makes no explicit commitment to locating natural kinds among the contents that explain experiences,

she thinks that duckhood of itself makes a phenomenal difference to visual experience (or at least it could, given her arguments). And I'd like in this section only to argue that nothing in her argument decides between the subject's having a recognitional capacity for ducks and a recognitional capacity for some cluster of visible properties.

To begin, it's rather implausible that the visual system can pick out ducks. One way to see this is to consider on the one hand what it takes to qualify as a duck, and on the other hand the nature of the visual sensors. I don't know what makes something a duck, and there is really no resounding agreement about what counts for membership in a kind, but the sorts of things that get suggested as candidates—born of ducks, having duck genes, etc.—are none of them directly visible properties. This is not to say that the containing system cannot be interested in ducks, and indeed it might very well use certain visible features of ducks to pick out ducks in some indirect sense. After all, the containing system can come to believe that something is a duck by looking. But that needn't require the *visual* system ever to token *duck*. It need only token some cluster of visible properties that reliably covary with ducks. If the containing system is disposed to associate that cluster with duckhood, then the tokening of that cluster will dispose it to believe that there is a duck present.

If this is what's going on when people look at ducks, then it's no wonder why, if Siegel is right, even an expert will be fooled (perceptually) every time she sees a suitably designed decoy duck. And it's important to acknowledge that Siegel needs to explain this fact. She concedes that an arborist will token *elm* upon viewing a suitably designed hologram of an elm.[14] And we can ask why this might be. And I think it is precisely because the visual system is picking out a cluster of visible properties rather

[14]I take it that she uses the example of the hologram because holograms reproduce depth phenomenology, even under translations in space.

than duckhood (or elmhood) itself, which, we've suggested, is not itself directly visible. In fact, an expert will experience the (perhaps faint) feeling of recognition even when viewing cruder depictions of ducks. No such depiction will fool the expert in the doxastic sense, but I'd wager it would take very little to provoke at least some degree of the recognition felt when viewing actual ducks. And what the cruder depictions lack (visually), compared with actual ducks and perfectly designed robot ducks (say), are things like viewability from all angles, verisimilar motion, precise shape, etc.[15] And these are all things that a non-duck could possess. What could it be about actual ducks, over and above these features, that the visual system is hooking onto? If we can't find any such things, then I think Siegel at least owes us an argument for the claim that what gets picked out are ducks rather than certain visible properties associated with ducks.

Siegel could certainly agree with everything I've said here and point out simply that vision can token a representation it really has no business tokening in any environment where both ducks and decoy ducks reside. After all, vision can be wrong about decoy ducks. That ducks and decoy ducks are phenomenally indistinguishable does not of itself settle the question whether the visual system can token *duck* or is limited to tokening *duck-associated property cluster*. We should want vision to distinguish between ducks and decoy ducks, but of course in any ordinary environment, it wouldn't have to. So, why should I care whether the visual system tokens *duck* instead of some cluster of visible properties? It is for the reasons I touched on at 2.2.5: the token *duck* cannot properly explain duck-seeing experiences, because it is not experientially distinguishable from a token involving only a cluster of visible properties.

[15] There are certainly features accessible by other modalities as well.

Say that someone is an expert, and she is looking at a decoy duck yet tokening *duck*. Now consider the same expert looking at a duck and tokening *duck*. The experiences are the same, and the contents are the same, and so the *l*-supervenience of contents on experiences is preserved. But there is something unsatisfying about explaining the experience of the decoy duck in terms of the content *duck*. Yes, *duck* is the token, and so we should expect the experience to be a certain way. But why is *duck* tokened? It is precisely because vision cannot distinguish between the duck and the decoy duck. But it pretends to: it can token a representation that reaches beyond what is visible, but it cannot apply that token reliably. Thus, though it uses a token that promises to make the distinction between ducks and decoy ducks, it cannot make good on that promise. And this is why the explanation of the decoy experience in terms of that token seems unsatisfying. In fact, it likewise makes the explanation of the duck experience in terms of that token seem unsatisfying.

Consider an impoverished class of visible properties associated with ducks, and imagine those properties instantiated in a decoy duck. Now, add successively to this class of properties various of the other visible properties of ducks. As you get closer and closer to a decoy duck that will fool an expert every time, you reach a point where the only difference between the decoy and a real duck is, say, realistic motion. At this point, you can say that the difference between an experience of a duck and an experience of this decoy is explained by the lack of realistic motion properties in the decoy. Now add the realistic motion properties. Suddenly, according to Siegel, your visual system tokens *duck* during either experience, every time. It is not simply making a perceptual error; rather, it is attempting to wield a representational token it cannot handle. It is getting all the visible properties right, but it reaches beyond

those properties and says something of the decoy that is altogether hidden from view. The difference between the *duck* token and a token similar in every way but stripped of all the non-phenomenal baggage—call it '*duck**'—is a non-phenomenal difference. So, why, then, do we explain the phenomenal aspect of the experience in terms of the *duck* token? Is does nothing to modulate the experience to fit its meaning in any way distinct from the *duck** token. Thus, a significant portion of the token is doing no explanatory work. Moreover, there is explanatory work left undone if only *duck* is tokened, work that must be done by additional visible properties that would be shared between any particular duck and its decoy counterpart. Because of this, I contend that the *duck* token cannot count among the contents that explain experience.[16]

There is really nothing I can say to someone who nonetheless insists that the relevant token is *duck*, and that it is the only game in town, so if anything is to do the explaining, it has to be that token. We might say, for example, that a token can mean *duck* as long as it is acquired in an environment with ducks and no decoy ducks. If the meanings of our representational tokens are given this way, then perhaps there are decoy-viewing situations in which the *duck* token is all we've got. But even in such cases, I'd argue that what explains the experience is, at best, those parts of the token that make a phenomenal difference. This would go some way to explaining our feeling that various impoverished decoys and depictions nonetheless provoke some of the same feeling of recognition. Unless the *duck* token can be considered in parts like this, that observation is left unexplained. And if it can be considered in parts, then I really see no difference between saying that the experience tokens *duck* versus saying that it tokens *duck**. Additionally, as I've said, no particular experience is completely

[16] And in fact, neither can the *duck** token explain the relevant experience completely, since what is contained within *duck* is explanatorily incomplete, and *duck** is a proper subset of *duck*.

explained by the relevant parts of *duck*. In order to give a complete explanation of any particular experience, we need only parts of *duck* and many other kind-neutral sorts of visual properties that both particular ducks and particular decoys could share. The case for *duck* as the explanatory content here looks pretty grim.

But Siegel does not adopt any one theory of intentionality and in fact claims that her view is compatible with internalism about content [Siegel, 2010, p. 114]. It's being compatible depends, however, on the rejection of the claim that specifically natural kinds figure in the contents of experience. If this is true, then perhaps Siegel would be willing to accept everything I've said here. But I suspect she wouldn't. Were she to do so, she'd have to admit that whatever the cluster of visible properties I argue vision is hooking onto, it counts as what she calls a 'k-property'. K-properties are defined negatively, as any of those properties over and above things like "spatial properties, color, shape, motion, illumination and their coinstantiation in objects" [Siegel, 2010, p. 6]. I don't mean the cluster of properties to include anything over and above these, so it's unlikely that on my view, vision picks out a k-property. Perhaps a cluster of these properties counts as a k-property, but she doesn't say one way or the other.

One reason we might think that she'd reject my view is that she rejects a view she calls 'Shape-Gestalt' according to which vision picks out not ducks but duck-shaped gestalts.[Siegel, 2010, p. 110] She says she has no "knock-down argument" against that view, but she describes a potential problem for it. The problem for a view like that, she thinks, is defining the gestalt in such a way that it reliably picks out the duck-shapes and only the duck-shapes.[17] If the acquisition of such a

[17]Or perhaps more accurately, visible duck-properties. They needn't all be shapes.

gestalt produces a characteristic experience in the presence of duck-shapes and only duck-shapes, and if such experience is produced only after the acquisition of the gestalt (and not before, when the subject is still a naif), then perhaps it can do the job. And I see no reason why a suitable cluster of visible properties cannot play this role. After all, I've shown that there is nothing about the token *duck* that should allow it to latch onto duck-shapes any more effectively than does *duck**. The differences between ducks and decoy ducks are extra-visual differences, and so they cannot possibly make the difference between adequate and inadequate selectivity for duck-shapes.

2.4 Morrison and Perceptual Confidence

We've argued thus far, by small steps, that indeterminate contents cannot explain the phenomenology of perceptual experiences and so cannot be phenomenal contents. Additionally, we've argued that if contents are to explain experiences, those contents must be phenomenally bounded, and we've seen that this appears to rule out of consideration certain 'high-level' properties like natural kinds. Though indeterminacy generally and even perceptual indeterminacy are well-trod issues in philosophy, the particular form and purpose of my arguments here is relatively new.[18] Because of this, my arguments here have few direct opponents (or allies, for that matter) in the literature.

In a forthcoming paper titled "Perceptual Confidence" [Morrison, 2015], John Morrison gives what is perhaps the only well-developed case for indeterminate perceptual attribution currently in the literature. There is a complication, however, in

[18] As far as I can tell, the present form of the debate over indeterminate perceptual content—keyed as it is to issues surrounding the relationship between perceptual content and perceptual phenomenology—originated in [Block, 2010].

considering Morrison's view as against my own: he does not claim that the contents he's talking about are explanatory of experience. Thus, it could be that he thinks such contents are not phenomenal contents and therefore do not *l*-supervene on experiences. And I've argued that if contents are to explain experiences, they must *l*-supervene on those experiences. However, Morrison does say that he is drawn to a view he calls 'confidentialism' according to which changes in 'confidence' contents are always accompanied by experiential changes; on this view, then, confidence contents do indeed *l*-supervene on experiences. At least, the 'confidence' aspect of those contents *l*-supervenes on experiences.[19] If this is right, then the disputed contents are within the purview of my arguments, because my primary argument for the claim that indeterminate contents cannot explain experiences is that such contents cannot *l*-supervene on experiences. And if Morrison accepts confidentialism, then he believes that indeterminate contents *l*-supervene on experiences.

2.4.1 Perceptual Confidence

John Morrison argues for a view he calls 'perceptual confidence', hereafter 'PC'. In his words, PC is the view "that perceptual experiences assign degrees of confidence" [Morrison, 2015, p. 1]. And though he allows for various interpretations of 'assign', I will assume that the assignment of a degree of confidence involves indeterminate property attribution, in the following sense. Say that we have an object *o* and a shape *S*. Then below, [a] is determinate and [b] is indeterminate:

[19]Morrison says that even if this is true, it is possible that two experiences alike in phenomenology might yet represent different shapes, colors, illuminations, etc. He cites a Twin Earth example as an illustration of this possibility. As I've said, I am not defending global supervenience, and so I am not necessarily in disagreement with Morrison on this point.

[**a**]: o is S

[**b**]: o might be S

Now, we *can* say of [b] that it says determinately that o might be S. Certainly, saying that is not the same as saying that it is not the case that o might be S. So at least, we are deciding between these two alternatives. But the attribution of S to o is *not* determinate, and that is what matters for our purposes. Thus, we cannot identify indeterminate contents without first specifying which elements of the construction are the objects, which are the predicates, and which are the connectives. For, if we treat *might be* S as a predicate, [b] attributes that predicate determinately. If we treat S as the predicate, the attribution is indeterminate.

It's not entirely clear what it would mean for a perceptual content to involve a connective like 'might be'; even worse is the status of saying that o has S with $x\%$ confidence. For, at least in the former case, we can understand 'might be' as indicating possibility, and then there are two choices—possible and not possible. If percentages are brought into the picture, we are faced with the question what they are percentages of. The answer, of course, is 'confidence', but we don't know what that means. Morrison uses both metaphysical and epistemic language in his paper. The paper is about 'confidence'—which seems epistemic—but also uses constructions like 'could be', which seem metaphysical. And it's not crazy to think that something like metaphysical possibility might figure into our perceptual contents. Perhaps there are metaphysical relations between visible facing sides of objects-at-distances and object identities. And maybe perception is capturing those relations. But since Morrison wants us to understand perceptual confidence on the model of doxastic confidence,

it is more likely that he means something epistemic. Let's say for now that we can understand degrees of belief, and that perceptual confidence should be understood on that model, whatever it is.[20]

2.4.2 Perceptual Confidence and Doxastic Confidence

Morrison's argument involves the description of a number of cases where our doxastic confidence about some perceptual stimulus undergoes gradual change. Imagine, for example, moving toward an object, and as its visible features come into view, becoming gradually more certain about its identity. He asks, of such cases, how we are to explain this gradual change in doxastic confidence, provided that we completely trust our experience. Now, of course, he is not claiming that we always completely trust our experience. But if it is possible to completely trust one's experience in a case where doxastic confidence changes gradually, what explains this gradual change? He thinks the best explanation is a corresponding gradual change in perceptual confidence. And if perceptual confidence can change gradually, then PC is true.

Before I consider his argument, there is already a worry about his conclusion—i.e., about PC. He begins with a worry about explaining doxastic confidence and gives perceptual confidence as an explanation. What, then, is supposed to explain perceptual confidence? If we can appeal to perceptual confidence without further explanation, why can't we do the same with doxastic confidence? Here, Morrison may object that his call for explanation is not a general one in the sense that it would apply to any sort of indeterminate content whatever. He might say rather that on the

[20]I'm willing to accept that we need something like degrees of belief, but I still can't work out exactly what that comes to unless it concerns evidence.

assumption that there are both perceptual contents and doxastic contents, and that the one serves as a kind of ground for the other, indeterminacy in the one could very well be explained by indeterminacy in the other, but existing analyses have seemed to rule that alternative out of hand. And maybe it's time we explored that alternative. Fair enough. But I nonetheless think it worth asking how it is that perceptual contents come to be indeterminate, and any answer to that question might serve as an answer to the question about doxastic confidence. If, for example, perceptual confidence is explained by the system's gradual acquisition of evidence (whatever this would mean at the sub-perceptual level), then perhaps something similar is true in the case of doxastic confidence. Perhaps, for example, despite perception's always speaking determinately, its testimony fails to be sufficient evidence for some proposition we believe with $x\%$ confidence. Unfortunately, Morrison never gives an account of perceptual confidence itself, and so it's difficult to pursue this line much further. I do think, however, that most of Morrison's cases can be explained in something like the way just outlined. More on that later.

2.4.3 Assessing Morrison's Argument

To begin, I'm puzzled about why there is an explanatory worry here. Morrison already admits that perceptual and doxastic confidence can disagree. To explain that disagreement, he brings in the notion of 'trusting' one's perceptual experience [Morrison, 2015, p. 22]. In cases where perceptual and doxastic confidence disagree, the disagreement is explained by our having less than 100% trust in our perceptual experience. So, why can't this always be the explanation? Morrison maintains that we 'can always imagine what would happen if [we] were to completely trust [our]

experience' [Morrison, 2015, p. 14]. But there is no way to distinguish the notion of trust from doxastic confidence unless you already accept Morrison's view about perceptual confidence (i.e., that it can be other than 0% and 100%). Morrison wants us to imagine cases where two things are true: Our doxastic confidence about some perceptual stimulus changes gradually over time; and we 'completely trust' our perceptual experience. Now, on his view, these conditions are compatible because perceptual confidence can also change gradually. Thus, he can say that doxastic confidence changes from, say, 20% to 60% because perceptual confidence does the same, and we completely 'trust' our perceptual experience. However, if we deny that perceptual confidence can be anything other than 0% or 100%, the case is incoherent. Completely trusting an experience that always gives determinate contents—i.e., contents of either 0% or 100% confidence—just means having 100% doxastic confidence in whatever the contents of that experience. Completely trusting your experience is just believing whatever its content. So, it is not possible to have a gradual shift in doxastic confidence while completely trusting one's experience unless PC is true. But, in order for his argument to have any weight, he has to get his opponent to accept the coherence of the example without already accepting his view. For this reason, I think Morrison's argument fails. However, he discusses two cases I find genuinely troubling for my claim—experiences of the visual periphery and experiences of the visual threshold—though not exactly for the reasons he gives. I turn now to these cases.

First, it may be helpful in clarifying my position to show why it is that most of his examples fail to threaten my view. I've hinted at this above. The first class of examples points to uncertainty in the representation of properties that I argue cannot

figure in phenomenal contents. Morrison seems to think their figuring in contents is unproblematic, and the problem is explaining their being accompanied by shifting doxastic confidence. Well, why not just think that our shifting doxastic confidence is due to the fact that our perceptual experiences do not speak at all directly with respect to such contents? Take the example of the approaching person, believed with shifting confidence to be Isaac, someone we believe to be out of the country. Morrison says that the gradual increase in doxastic confidence as the person approaches is explained by the fact that perceptually, we become gradually more confident, and we trust our perceptual experience. I deny that vision ever represents particular people, and so I explain the shift in doxastic confidence by saying that vision is 'impoverished'—to use Morrison's term—with respect to contents like *Isaac*. What happens in such a case, I think, is that as the person approaches, (visual) evidence mounts that it is Isaac. And on my view, at the level of visual experience, all of the relevant evidence is determinate content.

The Visual Periphery

There is, however, a need for some further discussion here. There is a difference in what sorts of visual properties are determinate between an experience of a far-off person and one close-by. Likewise, a richer range of visual properties are determinate in the center of the visual field than in the periphery. Here, it may seem, is a knock against my view, since it seems that, since precise color and shape are not assessable at a distance or in the periphery, color and shape are indeterminate contents at a distance and in the periphery. My initial response to this charge is that it's at least possible to have determinate contents in the periphery, or attached to objects that are

far away. The fact that a determinate content in the periphery might fail to decide a center-field content needn't mean that either content is itself indeterminate.

But I'd like to argue further that peripheral contents are indeed determinate. There are actually three sorts of contents to consider here—confidence contents, intensionally-ranged contents, and extensionally-ranged contents. Let's begin with the confidence content, because it is Morrison's focus. I think, in fact, the confidence content is the worst off, for reasons I've discussed. Of the three, it is the only wholly non-attributional content. This means that it runs afoul of points made at 2.2.3 (pp. 37-38). There, I argued that because phenomenology is attributional, the contents that explain it must be as well. And notice that phenomenology is attributional *even in the periphery*. Covertly attend to your visual periphery. Now imagine some minor change to your experience there. Intuitively, the world now seems different. Perhaps you cannot tell exactly in what respect it seems different, but this needn't mean that you've tokened some indeterminate content. After all, *ex hypothesi*, you are able to distinguish between the two peripheral experiences, but if the contents are non-attributional, they do not speak about the world in ways that are distinguishable. Another way to see the point is to consider two identical experiences of the visual periphery. If the contents that attach to these experiences are non-attributional, then the world could very well be different across the two cases. But is this possibility left open by the experience? Once we've agreed that the experiences are phenomenally indistinguishable, I can't see how this can be.

Whether Morrison would accept this style of argument, I can't say. He seems open to the possibility that representational content fails to do justice to perceptual experience. For example, he considers the charge that perception cannot token

confidence contents because confidence contents reference an array of possibilities, and in perception, we see one thing at a time [Morrison, 2015, pp. 13-14]. His response is to point to the possibility of experiences that token ranged constructions and to say that in such cases, perception tokens multiple possibilities without presenting all of them as actual. Thus, if it is allowable in the ranged content case, it should be allowable in the confidence content case. I think he's probably right about the logical relationship between these two cases, but I think they are both impossible. Let's consider ranged contents.

An intensionally-ranged content is one that attributes a range, *qua* range, of possible values to some object. I've said already that there appears to be no counterpart in experience for the subtle meaning of a range of values. Put precisely, an intensionally-ranged content is a content cast at determinacy level $p+1$ but is attributional only at determinacy level p. It is not, strictly speaking, non-attributional, but what it attributes is a set of possibilities. But experience does not move you through (or present all at once) the possibilities at some higher level of determinacy—it presents a property at its local level and nothing more. The reasons offered in the confidence content case apply here as well. Even peripheral experiences can always be compared for sameness and difference. Say they are the same. If we attach intensionally-ranged contents to them, then the content speaks about the world in a way closed off by the experiences, because the experiences do not allow that the world differs between them. Now consider two peripheral experiences that are different. If we attach intensionally-ranged contents to these experiences, then, similarly, the content speaks about the world in a way closed off the experience, because the experiences say that the world is different in either case, while the contents leave it open that they

might be the same.

Now, if you don't like this argument, consider that, as concerns ranged contents, we have a choice between extensionally equivalent but intensionally distinct alternatives. For every experience e and every range of values at *p+1*, consistent with e, there is some content at p that contains them and can be determinately attributed in e. My claim is that the p content is a more faithful rendering of the phenomenology of the experience than the extensionally equivalent content at *p+1*. For example, we can agree that 'indeterminate', as concerns ranged contents, is context-relative. *Every* content is indeterminate relative to some standard. So, in order for a content to be presented *as* indeterminate in this way, it has, in some way or another, to make reference to the relevant contextual parameters. This is why I say that intensionally-ranged contents are *p+1* cast contents attributional at p. They are 'cast' at a level different from the level of their attribution, meaning that they contain some reference to these levels. If they don't, then I think we are no longer talking about an intensionally-ranged content but rather an extensionally-ranged one. Let us now consider the latter sort of case.[21]

I said above that the possibility of extensionally-ranged contents might make some trouble for my view. The worry concerns how it might be that objects can be presented as having phenomenally *different* but *compatible* properties. I've already sketched one possible response: to say simply that vision only ever represents at one

[21]It will turn out, in fact, that marginal cases like these are unimportant for my larger project, which is to give an account of the phenomenal contents of 'constant' perceptual representations. And in those cases, the indeterminacy proposal will be assessed as a potential savior for the constancy orthodoxy. In that argument, the possibility that contents in the visual periphery are indeterminate is irrelevant, because the puzzles I consider will all of them take place squarely in the fovea. I do, however, maintain that no phenomenal contents can be indeterminate in the way I've discussed here, and I am here attempting to show that this claim is motivated independently of my goals in the constancy debate.

level of the determinacy hierarchy, which would rule out the possibility of different-but-compatible contents. Problem solved. But there are reasons to think this implausible. For example, I am capable of making fewer visual discriminations in the periphery than I am in the fovea. This suggests that the representations do not exist at the same level of the determinacy hierarchy, because if they did, my discrimination capacity should remain unchanged.

Before giving a better response, let's remind ourselves precisely why all of this is supposed to threaten my view. My goal is to explain perceptual phenomenology in terms of phenomenal content. By my criterion of explanatory adequacy, the phenomenology and content of an experience must be a match with respect to their intentional aspects. I here face a dilemma: I can say either that perception does represent at multiple levels of determinacy or that it doesn't. If it doesn't, then as we've said, I am left to explain a change in discriminatory capacity. If it does, then it might seem I am committed to saying that there is some class of experiences to which different-but-compatible contents do justice.

In fact, I don't think I'm committed to saying that *different-but-compatible* is phenomenally presentable. That red and scarlet are in fact different-but-compatible needn't imply that phenomenal presentations of *red* and *scarlet* reveal this relation between them. Properties stand in all sorts of interesting relations with one another, and some of these are not phenomenally showable. I contend that different-but-compatible is one such relation. I have no knock-down argument here, but since there is no special reason to think otherwise, I don't need one. I think the reason it is not showable, in any case, is that compatibility concerns category memberships that cut across levels of determinacy. And if we understand levels of determinacy as contextual

parameters, we can say simply that *red* and *scarlet* are incommensurable contents because they come from different contexts. The present suggestion will emerge again in my discussion of color in chapter 4.

The Visual Threshold

Threshold experiences are in certain ways like peripheral experiences, because whatever is going on in these cases is difficult to settle by introspection. Maybe there is room in the phenomenology for more than just 'present' or 'absent'. I don't know. The right thing to say is that changes in even threshold-level phenomenology are always determinate content changes—present or absent. And then our 50% response rate—i.e., our 50% doxastic confidence that defines the threshold point—is explained some other way, perhaps by our serious lack of trust in our perceptual system, for example. I think that in such cases, responding to Morrison is not difficult—100% trust in our perception in such cases is impossible, and it'd be a real shame if it weren't, because it should very well be adaptive to lose trust in perception with respect to those stimuli of such low intensity that neural noise becomes a confound. But I think there is a way to make the objection stronger. Whether or not we 'trust' our perception, why is it not possible simply to read off what it is telling us about the world in threshold cases? After all, it's unlikely that subjects would report about threshold stimuli that they seem always present or not present. It is more likely that subjects are trying to 'trust' their perceptual experiences at 'face value', but are having some trouble knowing which alternatives are being presented to them. I'm guessing subjects in such cases will say things like, 'I wasn't sure whether it was there or not'. And whether they actually trust their perception or not, they are caught in

the two-alternative forced-choice paradigm and must press a button.

Now, it's quite possible that I am here conceding too much to my opponent, and it could be that we get a 50-50 response rate for threshold stimuli simply because the perceptual system is only 50% accurate for threshold stimuli, and subjects are simply reporting what they see with 50% accuracy. That would certainly explain the result. But let's imagine that subjects are having trouble deciding what is being presented to them perceptually. The minimal description of what's happening in cases like these is that the subject cannot decide, on the basis of her perceptual experience, whether the stimulus is present or absent. But again, as before, we face a choice between saying that the subject's experience indeterminately represents the presence of the stimulus and saying that the subject's experience simply doesn't speak about the presence of the stimulus. And as before, the fact that our experience is a certain way—irrespective of how it speaks about this or that state of affairs we might be interested in—suggests that the non-representation thesis is phenomenally adequate while the indeterminacy thesis is not. The 50-50 response rate for threshold stimuli is explained, then, by the fact that the subject's *experience* fails to pronounce on the task at hand while the experimental setup demands that the *subject* pronounce on the task at hand.

Summary

For these reasons, I do not think Morrison's arguments tell against my view. Thus far, I have argued from the desideratum of explaining perceptual experiences in terms of their contents to the claims that phenomenal contents 1) must be determinate; and 2) must involve properties that are perceptually bounded. Hereafter, I will set

aside (2) and focus on (1). The two claims are related, because phenomenally unbounded properties are likely to be indeterminately represented.[22] For example, say that perception cannot distinguish (phenomenally) between Isaac and imposter-Isaac. Presented with an Isaac-looking man, perception might represent him as indeterminately Isaac. My worry, therefore, about the presence of indistinguishable properties in the contents of perception is in part that such contents are bound to be sometimes indeterminate.[23] And I have been concerned to show that such contents cannot be indeterminate.

2.5 Wrap-up

In this chapter I have argued that the contents that explain experience must be determinate and phenomenally bounded. These contents not only explain experience but put in place a kind of epistemic foundation for higher, more abstract representations of the world. Representational systems like human brains are interested in an array of properties, very many of which do not always enjoy determinate and phenomenally unique perceptual representation. The puzzle for the next chapter is to explain how it is that representation at the phenomenal level serves as a foundation for representations of the sort that truly matter for the representing organism, and how representations like these affect representations at the phenomenal level, if at all.

[22]Moreover, some indeterminate contents are phenomenally indistinguishable, as I argued extensively above in the case of color.

[23]Additionally, I am worried that experientially indistinguishable contents cannot explain experiences, as I have argued.

Chapter 3

Underdetermination

3.1 Constancy and Phenomenal Content

I argued in the previous chapter that phenomenal contents—those contents that explain the phenomenal character of our perceptual experiences—cannot be indeterminate and must be phenomenally bound. This places some constraints on the kinds of properties that can figure in phenomenal content. In the next two chapters, I make an extended argument for the claim that the properties involved in what are commonly called perceptual 'constancies' never figure in phenomenal content as such. For example, I will argue that surface color never figures in phenomenal content on its own as a feature value; it contributes to phenomenal experience only as part of a color product of specially conflated properties, a color product I will call a 'stack'. The problem of the stack is a sensory underdetermination problem. It is commonly thought that perceptual content is underdetermined by sensor input, and I will argue further that phenomenal content is underdetermined by sensor input. This means that phenomenal content 'overreaches' the sensor inputs to generate a more sophisticated

picture of the world and in this way overcome the problem of underdetermination. However, as I've argued, phenomenal content is limited in that it is comprised only of properties that are determinately attributed and phenomenally bounded. This means, will shall see, that there are special underdetermination problems—the problems of the 'stack'—that are not phenomenally resolvable. These problems are partially resolved, in a way that falls short of providing bounded, determinate representation of the relevant properties in phenomenal content.

In this chapter, I will begin by discussing the problem of sensory underdetermination in general and subsequently the more specific and more difficult problems of the stack, or 'stacking problems'. I will then discuss phenomenalism, a view according to which phenomenal content is not underdetermined by sensor input. I will then discuss the ways in which phenomenal content is indeed underdetermined by sensor input and so 'overreaches' the senses to represent important worldly properties, and I will appeal to some of the relevant empirical literature to give the view some experimental backing. I will finish by distinguishing between two sorts of overreach, only one of which can give partial solutions to stacking problems.

3.2 Sensory Underdetermination

3.2.1 The Problem in General

Underdetermination is a well-known problem in perception. In computational terms, we might say that some value x is underdetermined by input I (perhaps some informational set) if we cannot solve for x given only I. A value x (for some object feature, say) is sensorily underdetermined when I is the distribution of energy across

the sensors. The definition is loose, to be sure, in part because the computational language suggests that there is a function into which the inputs are fed, and if there is, then we can do much more with the inputs than we could without it, and so perhaps the most basic description of sensory underdetermination locates the inputs just outside the perceptual system, or maybe just inside, at the first level of activation, to be sure our inputs are not already function values. In any case, those details are not our concern here.

There are various reasons that early sensor activation might fail to give values for some variable of interest; truly, the reason in all cases is that the variable of interest is measured directly neither by the sensor nor by some subset of the whole sensory array.[1] But there are more informative ways of describing the problem in particular cases. It helps to consider candidate properties-to-be-detected. The easiest to imagine is a property that simply has no causal impact on the sensor at all. Energies at frequencies high and low can fall outside a sensor's detection range, and such values for such properties are of course sensorily underdetermined. Similarly, a sensor may simply lack the acuity to determine a feature value at a given level of determinacy, even if the feature causally impacts the sensor. Very high spatial frequencies, e.g., are not resolvable by the human retina (this varies by luminance contrast), and are therefore underdetermined by sensor activation. These problems are not solvable by the perceptual system or even anything computationally downstream. There are areas of the (human) brain that can token representations involving such properties, of course (since we can think about them), but the perceptual system cannot signal their local instantiation.

[1] I'm not here assuming that the values are contents, though it sounds that way.

Some properties are too 'high-level', we might say, for our sensors to detect reliably. A property might be too high-level if it has no regular sensory signature, or if what distinguishes it from other properties is something imperceptible. In the previous chapter, I argued that duckhood is not a phenomenal content for similar reasons: that it has no regular phenomenal signature, since ducks can look very different from each other, and what distinguishes a duck from a cleverly-designed decoy duck is something imperceptible. Ducks are not 'phenomenally distinguishable' from decoy ducks. The sensory and the phenomenal are distinct, to be sure, but the same sort of argument can be made to show that duckhood has no regular sensory signature. Here, we recognize that a perceptual system might very well 'overreach itself'—i.e., token representations involving contents it cannot determine using only sensory information. For example, the system might just token *duck* anytime it detects something with a duck-like sensory profile, and it might do a good job, too, as long as there are no non-ducks around with identical sensory profiles. I have argued that the visual system does not token *duck*, but the point is controversial. Siegel, as we've seen, argues that the visual system can token representations involving natural kinds as contents; we might therefore see her view as describing a perceptual solution to an underdetermination problem involving natural kinds.[2] And it is important to recognize that every underdetermination solution is an overreaching, even if reliably accurate.[3] If the system is to overreach like this, it is likely to do so according to rules, and these rules can be implicit and built into the structure of the perceptual system, some of them high and some low—some even in the sense organs themselves, perhaps.[4]

[2]It occurs to me that even in the 'unsolvable' situations I describe, the perceptual *could* very well token the forbidden representations (prompted by what, I don't know), but I doubt any actual systems would have such foolish design.

[3]I ignore questions about whether such overreaching is 'justified' in an epistemic sense.

[4]Consider the pupil as an aperture, for example. I don't want to consider yet whether pupil

3.2.2 Stacking Problems

Underdetermination in high-level properties is a variety of conflation problem—since ducks and decoy ducks generate the same energy distributions across the sensors, the sensors conflate decoy duck surfaces with duck surfaces. But there are more insidious conflation problems, ones not solvable by, e.g., playing the odds with the local kind populations. These tougher problems I will call 'stacking problems', because they involve *co-local* conflation—i.e., conflation by the sensors of two properties represented in the same place at the same time. Color perception provides what is perhaps the cleanest example, one I will discuss in detail in the next chapter. It is well-known that the color signal—i.e., the composition of light that reaches the retina—is determined at some retinal location l by both the spectral reflectance of the surface s at l and the composition of the illuminant incident to s. Variations in either property cause precisely the same sorts of changes to the color signal at l. As such, we cannot distinguish between surface changes and illumination changes by the color signal alone. In the color signal, surface reflectance and illumination arrive 'stacked', so to speak, at each retinal location l. The problem is general, however, and applies beyond color. The only actual instances of stacking that I'm aware of, however, are the special difficulties that so-called 'constancy' mechanisms are supposed to solve.

Consider spatial extension. There are three dimensions to spatial extension, and when we consider the volume of space between the viewer and some seen object, we have also distance from the viewer to represent in addition to object size in three dimensions (or two, if you like). And it is well-known that the retinal size of an object varies inversely with distance from the viewer. Thus, retinal size—a measure of sensor

dilation overreaches to *accurate* representation or not; maybe it misrepresents for the greater good. If so, it's not a 'solution' in the sense I mean here.

activity—has simultaneously to measure planar size (the size of an object along the plane perpendicular to the viewing direction) and distance from the viewer, in the same place at the same time. If there is a change in one of the represented properties but not the other, the system has now to represent inconstancy in one property and constancy in another (if it is to represent accurately), while the sensor information, keyed equally to both properties, is itself inconstant. So it must manage very dynamic and irregular relationships between sensor information and object features: because distinct properties can impact the same sensors at the same time in the same way, the system cannot simply register a feature change anytime there is a change in the energy distribution across the sensors affected by that feature.

My interest in the problem of stacking concerns the possibility of phenomenal resolution. The problem, truly, is just sensor conflation at a location at a time and needn't have anything to do with phenomenal content. The problem at the sensors is problem enough for the perceptual system. And perceptual solutions to the problem needn't do any more than enable independent representation of the stacked properties—or perhaps, even more weakly, it need only enable independent discrimination of the relevant properties. One way it might do that is *via* independent representation, but I see no reason why that need be the only way, and even if it were, it needn't concern phenomenal representation. In any case, I am here giving an extended argument about phenomenal content, and so my interest in the matter concerns the possibility of phenomenal resolution of stacking problems. Full phenomenal resolution would involve nothing less than wholly independent phenomenal representation of stacked properties—put crudely, the system would need to divide the stack so the parts can be seen. This is a difficult problem—so difficult that in one sense it cannot be solved.

Or so I will argue. But before I do, I'd like first to discuss the ways in which problems of underdetermination can be solved. Before I can do that, however, I need a foil.

3.3 Phenomenalism

Faced with the problem of sensory underdetermination, one might adopt a sort of conciliatory epistemology according to which all of our 'high-level' object talk—which talk got us into trouble in the first place—is really just talk about sets of actual or possible experiences and their features. The properties with which we are acquainted, on this view, are not material objects in space, but properties of sense-data, the immediate and infallibly introspectible objects of our perceptual experiences.

Phenomenalism is motivated by a very thoroughgoing empiricism. It solves the problem of giving an empiricist justification of our belief in the external world and its objects by giving up the search for evidence and insisting instead that our beliefs aren't what we thought they were. If our beliefs are, at bottom, about aspects of our immediate experience, then empiricism is vindicated and we can all go home. Not surprisingly, this approach has struck many as altogether unsatisfactory. The metaphysics of phenomenalism is, indeed, quite revisionary, and we might object to it on those grounds alone. Perhaps worse is that it's unclear how phenomenalism can handle the problem of error—i.e., the problem of making possible representational inaccuracy.

The *non-phenomenalist* sense-data view is thought to do a good job of accounting for error. In fact, it is typically the 'Argument from Illusion' that motivates the view in the first place. That argument asks how it is that an object can look to have properties it doesn't have (illusion), and how one can have an experience

as of some-or-other instantiated properties when in fact there is nothing there at all (hallucination). It answers that in these cases, though the external objects in question either don't exist or don't have the properties they seem to have, we directly perceive sense-data that do have such properties. I explain my illusion that a has P by identifying a sense-datum that both has P and is the immediate object of my illusory experience. I explain my hallucination as of a P by doing the same. The argument finishes with the claim that all we ever directly perceive are sense-data. On this view, there is a question to be answered about the epistemic relation between sense-data and their (hypothesized) causes, but assuming we can give an answer, we have dealt with the problem of error.

Phenomenalism, however, might appear to lose this advantage. The phenomenalist does not suppose that sense-data are caused by external objects. She says instead that objects (and certain properties) are constructions of sense-data. But the appeal to external causes of sense-data is an essential part of the sense-data theorist's Argument from Illusion, because it is in the relation between sense-data and their hypothesized causes that she defines hallucination and illusion. The phenomenalist can make no such appeal. Whether the phenomenalist might try to solve this problem in one way or another is not our concern. However, there is a lesson here about underdetermination and overreaching.

The sense-data theorist's Argument from Illusion applies more widely than I described above, in fact. Its conclusion is designed to handle, in one sense, varieties of the stacking problem I introduced above. Consider the stacking problem of spatial extension (or size and distance, or 3D shape—they are all the same problem). This is where, in the sense-data literature, we are asked to imagine viewing a coin from

various angles.[5] We want to say that in one sense, from certain angles, the coin looks round, and from other angles, oblong. But we don't suppose the coin actually changes in shape with the viewing angle. The sense-data view is supposed to solve the puzzle by identifying sense-data that actually have oblong or circular shapes.[6] Collections of these sense-data are caused by the taking up of several perspectives on an object with a stable three-dimensional shape; thus, to experience an oblong sense-datum is to the see coin in depth.[7] I think that this account underdescribes the experience here quite a bit, and I will say why in the next section. For the moment, I want only to note that this account does have the general structure we'd need to account for depth perception.

Phenomenalism, I think, does not. For to give an account of our representing objects in depth is to give an account of our representing an item's distance from the viewer (say, for example, that the relevant item is some undetached region of an object). And to admit of such contents is to admit that we represent the world outside as such, and this is to reject phenomenalism. I will argue below that depth is a phenomenal content—a form of overreach, to be sure, but a phenomenal content nonetheless. And this means that depth is given in the experience. It is not an inference from multiple perspectives, and it is not 'built up' out of variously-shaped sense-data. A single view of an object contains a phenomenally salient representation of depth, and so to construct a world out of the elements of experience is to construct

[5]See, e.g., [Ayer, 1963, pp. 3-5]

[6]Other of the stacking problems are handled similarly—like surface color in changing illumination contexts, e.g.

[7]Or maybe the case is said to be about shape constancy. But the problem of 3D shape is really just a variation of the problem of planar size and depth, since shape is the complete set of measurements across the surfaces of an object (and invariant to scale). This is not to say that the brain doesn't have special resources for processing shapes, as apart from its size and distance processing. Maybe it does.

a world in depth. And this the strict phenomenalist cannot accept, because only a world external to the perceiver can have depth.[8] If this is an accurate account of the phenomenology of experience, then Hume and Berkeley are mistaken: there are perceptual representations of externality. They arrive packed into single perceptual tokens and are not complex inventions of the imagination or fictions collected into sets of perceptual tokens.

Similar points apply to other of the stacking problems, and this is in general because stacking problems concern worldly properties, ones that matter a great deal for our worldly projects (like not dying), but perception seems, or might seem (if we aren't careful) to give us nothing but basic, uninterpreted experience properties—things like planar shape, apparent size, apparent color, etc. If this is our view of perceptual experience, then it may seem reasonable for an empiricist to adopt phenomenalism's impoverished picture of the perceptual world. For, as I will argue, this is an impoverished picture of perceptual experience. It is a picture that, either by bad phenomenology or for reasons related to the larger epistemological project, sees perceptual experience as giving the perceiver an infallible perceptual foundation for world-construction—infallible because the foundation is composed of sense-data, about whose properties we cannot be mistaken. From this starting point, it would be damning to concede that perception regularly 'overreaches' in ways that modulate perceptual experience, ways that make possible the worldly sorts of contents that matter most for creatures living in the world. But perception does overreach like this—it tells us things about the world that, on the phenomenalist's strict epistemology, it has no business telling us, things whose truth values are not settled by the experience

[8] Actually, it doesn't really matter if the world really has depth. It's enough that our experiences say that it does.

itself. It cannot, therefore serve the very specific epistemic role it is meant to serve in the theory. Moreover, phenomenalism is therefore false.

3.4 Phenomenal Content and Overreach

3.4.1 Sensory and Phenomenal

In Chapter 2, I argued that phenomenal contents must be determinate, and that the properties that figure in the phenomenal contents must be phenomenally bounded. Loosely, the constraints on phenomenal content that fall out of these conditions suggest that phenomenal contents are generally 'low-level'. The language of levels in psychology typically assumes that the lowest level is the sensory and that level measures are approximately distances from the sensory. And maybe this is measured in processing steps, or something like that. When I use the term here, I am very consciously envisioning a very tight connection between the sensory and the phenomenal.

To hammer an already tired example, an initially plausible reason that duckhood never figures in phenomenal content is that the essential properties of ducks—being born of ducks, having duck genes (say)—do not impact the sensors. True, there are links between some essential properties of ducks and physical properties (phenotypic ones) that impact the sensors, but these properties can turn up in decoy ducks, too. And if ducks and decoy ducks share all of their sensory properties, then those properties cannot comprise a sensory 'signature' for either (or perhaps it is a signature for a category that includes both). Note that this does not depend on the actual incidence of decoy ducks amongst local populations of real ducks. Even if, in the actual world,

there are no decoy ducks, and the sensory profile for ducks does therefore count as an actual-world signature, the property of duckhood and the sensory profile differ extensionally in possible worlds.[9] It is therefore seems that duckhood, though surely a representational content of some sort, is not a phenomenal content.

In rejecting phenomenalism, however, I made the claim that the phenomenalist imagines perceptual experience too basic, too uninterpreted, non-worldly—too merely sensory, we might say. It's as if the phenomenalist, concerned as she is to make perceptual experience an epistemic foundation, cannot bring herself to see anything in the content of experience beyond what sensor information can strictly determine. I do think that basic sensor information is indeed a major contributor to phenomenal content. Consider, for example, the specificity and particularity of phenomenal experience. Surely this is explained in large part by the contribution of sensory information to our perceptual experience. Even if, for example, our 'conceptual scheme' has some impact on 'what we see', conceptual abstractions can hardly account for this aspect of perceptual experience. Moreover, sensory information marks a kind of starting point for our knowledge of the world, since it is only through the sensors that the outside world is able to make any causal impact on our mental lives. Without the senses, and without an inherited sensory history, we'd have no access to the world at all. But for any particular perceptual experience, had by a particular subject, sensor information tells only part of the story. The rest has probably to do with our inherited sensory history, among other things, but I am not concerned to speculate about that. I want first simply to pump the intuition that the merely sensory underdescribes perceptual experience. In fact, the merely sensory *grossly* underdescribes perceptual

[9]It's worth pointing out that even young children have the intuition that natural kinds are extra-perceptual. See [Gelman and Markman, 1987].

experience. Among the likely phenomenal contents underdetermined by the sensors are surface relations, depth and texture, 3D shape, size relations, object motion, object individuation, contour grouping, certain shape Gestalts, certain visual phenomena like shadows (e.g.) and many others.[10] Some of these contents figure in solutions to stacking problems; others figure in more garden-variety underdetermination problems. I'll consider as examples the properties of visual grouping (a non-stacking problem) and depth (a stacking problem).

3.4.2 Grouping and Depth (and Ducks Again)

In Chapter 2, I argued that Siegel's view fails to explain why it is that duck experts experience duck-recognition phenomenology in the presence of so many different non-ducks. If it is that easy to 'fool' the expert, maybe vision is hooking onto some property other than duckhood. This would, after all, make veridical a greater share of the expert's experiences. But this argument threatens to rule out too much. It threatens to rule out any property that can be reliably faked. I will focus on grouping and depth in particular.

I should start by convincing the reader that grouping and depth are both distinct in a certain way from other visible properties and important to preserve as contents of visual experiences. If you consider the case of color, you'll see that the argument made above against *duck* fails pretty straightforwardly. After all, any mock-up or decoy of a color representation will actually have whatever color it is depicted as having.[11] The same goes for texture, and perhaps for 2D shape as well.

[10] Aspects of my view place a 'ceiling' here, however. Quite plainly, the added contents tend towards 'higher-level' sorts of properties, but as I've argued, properties at too high a level simply lack phenomenal signatures or perhaps cannot figure determinately in experience. If so, they cannot be phenomenal contents.

[11] I am here setting aside the issue of the conflation of surface color and illuminant. Let's say for

These properties cannot be faked. But grouping and depth can be faked.

Figure 3.1: An example of ambiguous grouping.

In the figure above, the grouping is ambiguous between two interpretations. On one reading, the fleshy region at bottom-center is the man's arm. On another reading, it belongs to the woman's arm (it is in fact her arm). Notice that your experience of the image is different depending on which reading you choose. And it is easy to modulate your visual system's perception of grouping, even under optimal perceptual conditions. Similar points can be made about depth, and we can see in the above image that depth is likewise being faked.[12] And their being 'faked' is a fact about the experience, not what we are disposed to believe because of it. I, of course, do not believe that there are 3D human limbs in front of me, but it still looks like there are. And this suggests that grouping and depth are essential parts of visual

the time being that colors are just mixtures of light that hit the retina. I'll deal with perceptual constancies in detail in the next chapter.

[12]It is also possible to fake depth using 3D objects—you set them up so that the visual system tokens the wrong depth relations. I remember one such vivid experience looking at the netting between my seat and the arena floor at Medieval Times in Schaumburg, IL. I got to where I could make a Gestalt-switch between two depth interpretations.

experience, because I take it that I am representing, phenomenally, that there are certain groupings rather than others, and that I am looking at 3D objects.

Now, at this point, you might wonder why any of this is a problem. Grouping and depth have a certain phenomenal signature, and the fact that in this image that signature is provoked by illusion means only that the visual system can sometimes screw up and token depth and grouping properties in their absence. But recall the structure of the argument against *duck*. Siegel argues precisely that *duck* has a certain phenomenal signature, and I pointed to cases where that signature is reliably provoked by decoys. In the case of grouping and depth, you can likewise fool the visual system reliably in this way—in the case of 2D depiction of depth, e.g., the presence of a 2D array of shapes produces a reliable experience of depth. And notice that in all of these cases, the perceptual conditions can be as optimal as you like. So, what basis have we for saying on the one hand that vision does not represent ducks but instead represents a property shared by ducks and duck-decoys, and on the other hand saying that vision does represent groupings and depth, despite the fact that some humbler matrix of properties can produce the very same experiences? If we cannot distinguish between these two sorts of cases, then perhaps we are stuck saying that even when viewing 3D objects, we represent phenomenally only things like color, 2D shape, texture, etc., and that what we thought was depth and grouping phenomenology is really just the phenomenology associated with whatever matrix of properties both experiences actually share. In the duck case, we said that vision cannot distinguish between ducks and decoy ducks. In the present case, we might say similarly that vision cannot distinguish between actual depth and *depth**—that property we build into our 2D depictions of depth.

The solution to this problem, I think, helps us to better understand why we've ruled out *duck*. Siegel's (extrapolated) description of the duck case is that when an expert sees a decoy duck, she tokens *duck* and is mistaken. The Siegelian description of the grouping and depth cases is that when an 'expert'—i.e., someone who can group and represent depth—sees 2D depictions of depth, or tricky bistable grouping stimuli, she tokens *depth* and *this grouping* (or whatever) and is likewise mistaken. So, what's the difference between the two cases? I think that in the grouping and depth cases, the illusion interpretation is far more plausible than in the duck case. True, in either case, if the token is what Siegel says it is, then we have illusion. And if the token is what my argument suggests it is, then neither case is an illusion. But the illusion interpretation is so much more plausible in the present case than in the duck case. It's tough to say why exactly, but perhaps it's that the visual system has an urgent need to distinguish between depth and fake depth (since fake depth is not particularly unusual, even in the natural world) and no such need in the case of ducks and decoys (since true decoys are rare or nonexistent).

Or consider that, in nearly any actual circumstance, a *depth* token can be borne out (or showed false) by further perceptual exploration in a way that a *duck* token cannot. If you are presented with a 2D depiction of depth, and you view it from an oblique angle, you will lose the depth illusion.[13] Similarly, if you manipulate the objects in a scene, you can remove grouping illusions.[14] But what analog has the duck case? Investigate the decoy all you like, *ex hypothesi* it actually shares all of

[13]This is not true of holograms, but holograms will eventually fail some perceptual investigation. At the very least, you can disrupt a hologram using occluding surfaces in a way that you cannot with real 3D objects.

[14]It is of course possible to imagine depth illusions that cannot be destroyed by perceptual exploration. Thus, the possibility of destruction by exploration is not a necessary condition for illusion. It's more like a helpful diagnostic in most cases.

its perceptible properties with ducks. Because of this, we feel sort of guilty accusing vision of mucking up the duck case. What chance has vision to distinguish between visually identical items? Because it has no such chance, we are more inclined to say of this case that vision does get it right. And if vision gets it right, even in the decoy case, then it didn't token *duck*.

All of this is supposed to show that, despite our ruling out of phenomenal content certain properties that are not strictly determined by sensor information, there are properties underdetermined by sensor information that figure in phenomenal content. This means that phenomenal content overreaches, and it is by overreaching that it solves a problem of sensory underdetermination. The case of grouping was offered as an example where the system achieves full phenomenal resolution of the underdetermination problem. This means that grouping cannot be a stacking problem, since stacking problems cannot be fully phenomenally resolved. I used here as an example of a stacking problem the case of depth, and if we understand depth as part of the problem of representing spatial extension in three dimensions, then to solve for depth is to solve a 'stacking problem', in my terminology (grouping is not a stacking issue, but involves underdetermination nonetheless).[15]

And this means that the connection between sensory and phenomenal is looser than we thought. The best we can say about the connection between the sensory and the phenomenal is that one way a property might figure in the phenomenal content is by being strictly determined by sensor information. But it's not the only way, and it's

[15]There are very characteristic ways in which the perceptual system solves problems of underdetermination. Typically, the system estimates or otherwise assumes a needed value, likely on the basis of historical perceptual regularities or some such thing. I will not explore computational solutions in any detail, but it is worth noting that they all count as forms of sensory overreach because they compute on values underdetermined by the sensors. This makes them vulnerable in one sense (they've reached beyond where they strictly ought to), but their risky behavior makes intelligible the perceptible world.

not likely even the primary way.

3.4.3 Energy Map and Worldmaker Contents

There are, I contend, two distinct ways that a neural output can figure in phenomenal content. I will henceforth treat this distinction as a distinction between types of contents. The first sort I will call *energy map* contents, because they present (more or less) the distribution of energy across the sensors, and the second sort *worldmaker* contents, because they involve the properties we think of as fundamentally external, worldly and important. A crude way to understand the distinction in the visual case is to imagine a snapshot of your visual experience, like a screen-grab of your visual field. You can reproduce this image, one might think, by filling in point intensities at each location on the 'canvas', thereby generating an image that, if you were to look at it from the proper distance, would provoke in you precisely the same visual experience you were having at the time of the snapshot. Energy map contents are specified in the filling instructions used to produce the image. Worldmaker contents are captured in all those additional aspects of the experience provoked when you look at it. Consider that although there are no actual shapes in the image (just sets of point-intensities), you will represent shapes if you look at it. Let's say that in the brain, sensor inputs (or some selection of them) are sent to some higher area and computations there decide that there is a square object in the visual field. This result is then fed back to, say, early cortex, and only then do you have the phenomenal experience as of squarehood, even though all of the relevant point intensities were in place beforehand. According to this view, the total phenomenal content is constructed using two pathways, one from sensors to early cortical areas (say), and another from

the sensors all the way up to higher areas and back down again.

This means, however, that both sorts of contents will likely involve overreach, and both pathways are likely to involve both sorts of contents. Figure 3.2 shows a crude diagram of the view I'm suggesting, and it's clear from that diagram that there is some space between the sensors and 'seat of phenomenal experience', assuming there is such a thing.[16] We know, in fact, that there is a great deal of processing that happens between the sensors and early visual cortex (one candidate for the phenomenal seat). And this means both that there are likely some worldmaker contents that originate in the bottom-up pathway and that there is overreach in the energy map. In fact, both of the worldmaker properties I've discussed thus far—grouping and depth—may very well originate in the bottom-up pathway or in both.[17]

But I do think that a distinction between the two sorts of contents manifests phenomenally. You can appreciate the difference by considering bistable percepts, like the Necker cube:

Most people can easily train themselves to switch between two depth interpretations, and with a little effort, it is possible even to see the lines as a two-dimensional pattern. The signature of the energy map is the hard, dark imposition of the lines,

[16][Mumford, 1991] has a view like this, though his phenomenal 'seat' is the thalamus. [Lamme and Roelfsema, 2000] have a similar view with V1 as the phenomenal seat.

[17]Contents may originate in both pathways but according to different computations. It is thought that figure-ground segmentation depends on both 'bottom-up' computations that rely entirely on implicit rules about the grouping of stimuli, and 'top-down' computations that involve object categories and identity. We might say of the first sort of computation that it segments 'blindly' in comparison with computations of the second sort, although both sorts of computations may fall short of involving explicit representations of object features. That is, neither may involve computations that involve information at the level of perceptual content. If this is right, then perhaps the distinction between the two is merely directional. The vision science literature sometimes makes a distinction between bottom-up and top-down influences on perceptual experience by saying that the first sort are 'data-driven' and the second-sort 'concept-driven', but it's not at all clear whether the top-down computations of figure-ground segmentation involve concepts in the philosophical sense. I am not here adopting a view according to which top-down computations begin at the cognitive level of processing, and I'm doubtful that any perceptual computation involves contents in the philosophical sense.

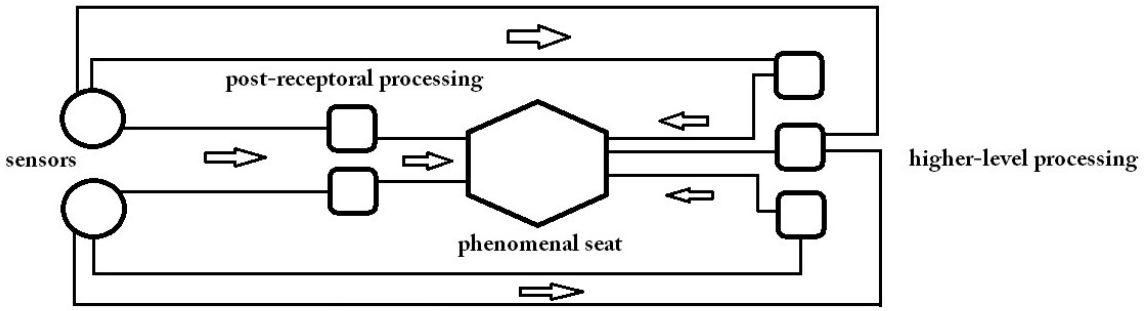

Figure 3.2: A possible schematic for two processing pathways.

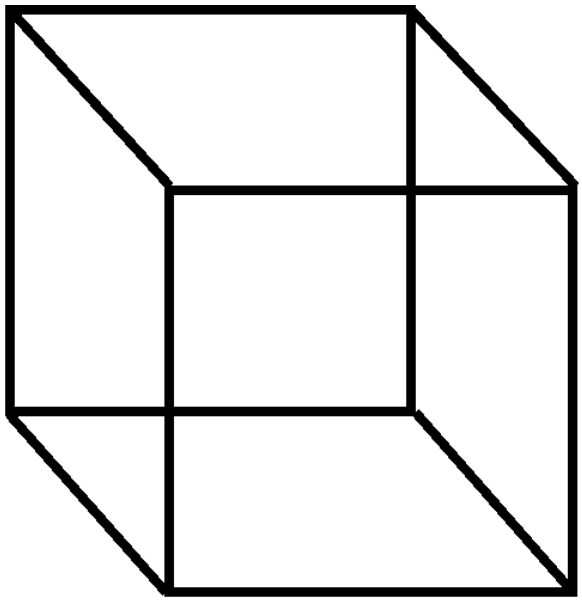

Figure 3.3: A Necker cube.

visually striking and irresistible. The signature of the worldmaker is the faint, almost kinesthetic feeling you get with each depth interpretation, especially as these compare with the relatively bland experience of the two-dimensional pattern. Additionally, the energy map is in some sense, at least in the visual case, 'saturated'—e.g., we could give a point intensity for every location in the visual field—and this means that worldmaker contents cannot occupy the same visual space as energy map contents, at least not in the same way. We might say that worldmaker contents are 'overlaid' or

'projected onto' the energy map, or even indexed to objects or features, rather than to retinotopic coordinates. I'm using somewhat impressionistic language here, of course, but I imagine most people can pick out the difference on their own. In any case, I think it's quite clear that the difference is part of the experience, whatever else it is.

Furthermore, the difference is related to the ongoing discussion about underdetermination and phenomenal resolution by overreach. Unfortunately for us, it is likely that both the energy map and the worldmaker involve overreach in the sense that their contents are neither of them strictly determined by sensor input. Thus, we cannot distinguish between the two by saying that one overreaches while the other doesn't. However, all energy map contents retain the 'merely sensory' representational format—point-intensities at retinal locations.[18] Energy map contents read like a roll call of sensor activations, and overreach at this level does nothing more than modulate that activity (it remains to be seen why this would be representationally useful). Worldmaker overreach, however, is both underdetermined by sensor input (and by the energy map, for that matter) and is different in format, and so its representations figure in experience in a quite different way. Detailed discussion of the matter will have to wait for the very specific examples of the next chapter.

For now, it remains to be seen whether this rather rough picture of neural connectivity and processing is at all empirically plausible. In the next section, I will try to give the picture some experimental grounding.

[18]I'm getting a bit chauvinistic here about vision, but the view is meant to apply more widely. I won't explore the matter, however.

3.4.4 The Role of Feedback in Awareness

There is a surge of interest in contemporary vision science concerning the role of feedback connections in visual awareness. It is well-known that higher areas of the visual system have extensive feedback connections with retinotopic cortex, but there is little widespread agreement about the role of these connections in vision, and of course the roles are likely to be several.[19] Important for our purposes are recent claims that feedback connections from areas like LOC to striate cortex are necessary for visual awareness.[20] More specifically important for our purposes is the possible necessity of such connections for certain kinds of visual awareness ('awareness' is a term used in the scientific literature that corresponds closely with what I'm calling 'phenomenal experience'). [Wokke et al., 2012] found that visual experience of perceptual completion can be altered by disrupting feedback from LOC to V1-V2. Subjects whose feedback connections were disrupted failed to see illusory contours in Kanizsa-type stimuli (figure below). Minimally, this result suggests two things: 1) LOC is active in the generation of perceptual experience; and 2) it cannot be active unless its outputs are fed back to V1-V2. It is plausible further to suppose that the role of LOC in visual experience is in generating what I am calling 'worldmaker' content.

The present study focused on illusory contours, but in order that they are seen as the contours of a square, LOC must generate a representation of squarehood, and its feedback to V1-V2 must modulate the input representation, at least in part, according

[19]Feedback could modulate activity in retinotopic cortex by either inhibiting or facilitating, sharpening, etc. See [Murray et al., 2002] for discussion. It might also recruit neural populations that haven't been active, or even superimpose novel contents onto visual awareness. See [Wokke et al., 2012] and [Williams et al., 2008].

[20]See, e.g., [Pascual-Leone and Walsh, 2001] for awareness of motion.

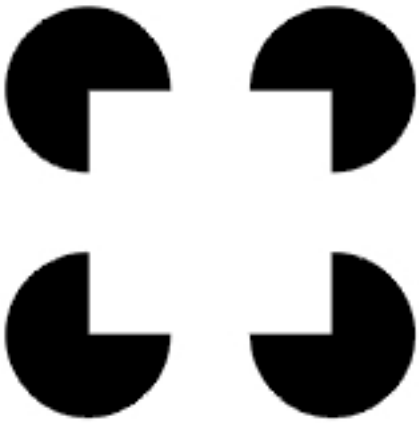

Figure 3.4: Kanizsa-type figure.

to this shape representation. I am suggesting, then, that shape-related phenomenal content makes its way into perceptual experience via these feedback connections.[21] So, on my view, the sensor input to V1-V2 gives the contours, and feedback from LOC (or somewhere like it) enables the subject to see collections of contours as collected—i.e., as the outlines of shapes or shaped objects.[22]

Feedback connections play a similar role in so-called 'predictive coding' theories [23]. Reaching back at least as far as Helmholtz is the idea that vision is 'hypothesis-driven'. Though the notion has not always enjoyed widespread popularity, there is a kind of renaissance currently underway, and at the center of this buzz are predictive coding theories.[24] Precursors of these contemporary theories tended to focus on the

[21]The visual system is not likely to have a mechanism for the express purpose of generating illusions. Perhaps the mechanism is used to overcome the effect of occlusion, and the present result is an artifact.

[22]Possible evidence in [Mevo et al., 2013] and [Pascual-Leone and Walsh, 2001].

[23]For well-developed statements of the theory, see, e.g., [Rao and Ballard, 1999], [Lee and Mumford, 2003], and [Friston, 2005]. Philosophers are now paying attention as well; see [Clark, 2013] for a discussion of the potential consequences of predictive coding theories for philosophy of mind.

[24]One might suppose that 'hypothesis-driven' would strike some as sounding too 'discursive' to describe something as sub-personal, as involuntary and, importantly, as *fast* as vision. Later carriers of the Helmholtzian mantle use different language. For example, [Stoner and Albright, 1992] refer

problem of coding efficiency—it is plain that in an environment with considerable visual redundancy, e.g., a visual system can achieve a favorable balance between accuracy and coding efficiency (and therefore, perhaps, metabolic efficiency and processing speed, reaction time, etc.) if it codes the redundant aspects of the environment sparsely. Very early incarnations of the idea seized on spatial redundancy—repeated spatial patterns, e.g.[25] More interesting for our purposes is temporal redundancy. A visual system might sparsely code familiar shapes, e.g., in order to save processing resources for novel stimuli.

But sparse coding is only part of the story. In order to exploit temporal redundancy, the system must make use of a predictive mechanism, and what drives the mechanism cannot be a fully-coded sensor representation, otherwise the advantage is lost. The basic story is as follows. Some subset of the sensor input is sent, via low-resolution and high-speed (probably magnocellular) pathways, to higher areas of the brain, like the lateral occipital area (LOC), e.g. Higher areas then use this information to register predictions about the stimulus, and the predictions are sent, via feedback connections, to early visual cortex, say V1, where the input representation is modulated in some way by the feedback. The result is that the coding in V1 is 'simplified' accordingly with the predictions. Hypotheses along these lines are typically driven by the oft-repeated result that 'predictable' stimuli increase activity in higher areas like LOC and reduce it in V1 [Murray et al., 2002, Lerner et al., 2001, Caclin et al., 2012].

But it is not coding efficiency that interests me here, since reduced activity in V1 could very well have little to no effect on perceptual experience.[26] And we

to the 'most likely interpretation' of a visual signal, which language might appear easier to identify with some neural process.

[25]See [Barlow, 1959].

[26]One issue with the predictive coding theories as they stand is that perceptual experience is suffuse with sensory detail—the sorts of details you'd expect to come from the sensors and show up

might ask, similarly, whether the predictive aspect of the theory has any role in my argument. I'll say that it is not altogether clear what feedback's being a 'prediction' amounts to in the first place. If feedback modulates V1 directly, then what we get is more than a prediction—it is a pronouncement. If, for example, distinct higher-areas give predictions about the same stimulus feature, and the predictions disagree, what happens in experience? I have argued that experience cannot reflect probabilistic outputs as such. All such outputs must pass through a decoding mechanism that selects a determinate answer. Perception might be wrong about a stimulus, but it cannot remain undecided in its effect on experience. Thus, any feedback that modulates visual experience is really, after all, just adding to or altering phenomenal content. And if Wokke *et al* are correct, at least some feedback does modulate visual experience, and it does that by modulating activity in V1-V2.

Moreover, I think that the motivation for calling the mechanism 'predictive' is simply that, in my terminology, it overreaches. That perception is 'hypothesis-driven' is supposed to be a surprising insight because it calls into question the 'purity' of the bottom-up picture of sensory representation. That general picture plays a central role in the phenomenalist epistemology described above, and it depends on viewing the contents of perceptual experience as infallible because metaphysically unsophisticated. But the contents of perceptual experience are richer than that picture suggests. And this means that what is distinct about the new orthodoxy is that it sees perceptual

in V1—like retinotopy, e.g. But if predictive feedback is supposed to 'simplify' V1 representations, why is it that everything that feeds into V1 appears to remain part of the experience? We might think that the feedback enables a kind of 'code-sharing' between V1 and higher areas, but then it's no longer clear why this strategy is any more efficient from a coding perspective, since in order to do the coding work V1 does, LOC, e.g., would have to code for retinotopy and everything else. LOC cannot tell V1 simply that there is a square present; it has to say exactly what the square looks like and where it is located, since our experiences of squares are always highly specific like this.

experience itself as a kind of overreach.[27] If this is true, then the fact that a top-down mechanism uses a crude selection of sensor inputs (for speed) and then feeds its 'predictions' back down to those areas of early vision thought to be responsible for 'awareness' is really no different than saying that there are top-down phenomenal contents, all of which involve overreach (worldmaker contents, plausibly). Perhaps the biological purpose of the mechanism is metabolic efficiency, but even if this weren't an issue, the system would need to overreach sensor representations anyway, because without overreaching, it suffers the ills of sensory underdetermination, including stacking problems. Unless a perceptual system can overcome these difficulties, it cannot provide a useful picture of the perceptible world.

3.4.5 Two Types of Overreach

As I've said, my view makes a distinction between two sorts of phenomenal contents, and the distinction is loosely correlated with origin points: energy map contents come from the sensors, and worldmaker contents from 'above', though not necessarily post-phenomenal seat (striate cortex, say), as I've said. My terminology picks out roughly the same distinction as 'bottom-up' and 'top-down', but I hesitate to say 'top-down', since I don't imagine that much of what I'm talking about originates in particularly 'cognitive' areas of the brain, which we might view as comprising the top of the processing chain. I do suspect that phenomenal content is 'impenetrable' to cognitive influence, in the following sense: cognitive contents cannot figure in

[27] Calling the view 'new' is misleading; there are many precursors to the view that make similar claims about the contents of experience, including Helmholtz, as I've said. Additionally, Bruner and Postman (See, e.g., [Bruner et al., 1949]) and others working in the 'New Look' paradigm uncovered a host of perceptual phenomena that can be viewed as mechanisms of overreach. The predictive coding theory is one modern incarnation of this general impulse.

phenomenal content *as such*. I have no argument for this claim, other than that I can't see how a cognitive content could possibly turn up *as that content* in phenomenal experience.[28] It is in my favor, then, that all of the alleged cases of cognitive penetration in the literature are cases where some cognitive state affects existing phenomenal contents (in a way that makes sense of the content of the cognitive state) rather than introducing novel contents.

An example discussed by Fiona Macpherson [Macpherson, 2012] involves the cognitive modulation of color contents. Subjects see heart-shaped cutouts as redder than they are (or anyway, their matching behavior suggests it). The effect is meaningful because subjects know hearts to be red, and this makes sense of their seeing heart-shaped objects as redder than they are. But the bit of knowledge—that hearts are red—is not itself a phenomenal content. That fact has no regular phenomenal signature. In this case, there is top-down modulation of phenomenal content, but the modulation seems not to be associated with a worldmaker content. Rather, it is a modulation to the energy map.

This means that there are both energy map and worldmaker contents (or content-modulations) coming from both directions, all of it overreach except the initial distribution of energy across the sensors. Moreover, it might be difficult in some cases to pick apart an energy map modulation from a worldmaker modulation, because the two might be related. If, for example, a change in apparent lightness makes an object more phenomenally salient—the phenomenon of 'pop-out', perhaps—or changes the depth or grouping contents associated with it, which sort of change is it? Both?

[28] One might think that certain hallucinatory states could be like this. Say that you believe that there is a gun on the table (out of hysterical fear, say), and this causes you to 'see' one there. I still think what you see is a perceptual content related to the belief, but maybe someone could take this line.

[Hupe et al., 1998] found that cortical feedback improves figure-ground segmentation, but it's not clear whether one or both sorts of change happened in experience. Peter Tse [Tse, 2005] describes an illusion where lightness and depth ordering interact.

Given these obstacles to sustaining the distinction between these sorts of contents, why do I insist on it? I insist on it because I think that only worldmaker contents have the potential to resolve stacking problems phenomenally. Remember that overreach, though it sounds epistemically dangerous, is a good thing. We need overreach; without it, the world would be largely imperceptible. And in fact, I am not advocating for a view according to which we must 'misrepresent for the greater good', or some such thing. I think that overreach brings us closer to accurate representation, and perhaps even justified representation (on a reliabilist account, say), though I'm not interested in this aspect of the view.

And indeed, I think that even energy map overreach could increase accuracy in some cases, at least indirectly. I mentioned above a possible role for lightness contrast in figure-ground segmentation. Consider that the Gelb effect (Figure 3.5) generates an illusion, but it might very well help the visual system in grouping tasks, and if it does, I think it'd be right to say not only that our figure-ground performance is improved, but that our figure-ground representations are more accurate.

Figure 3.5: The Gelb effect.

However, I do think that grouping contents are worldmaker contents, and so it's

unclear whether in this case the energy map modulation has itself made the experience more accurate or not. In any case, I don't think that energy map modulations can possibly solve stacking problems. I will make that argument in detail in chapter 5, but briefly, the reason is simply that stacking problems are the result of a special kind of sensory conflation—one where the conflated properties are phenomenally 'stacked'—and energy map modulations can do no more than give a new stack. The conflation remains.

3.5 Looking Ahead

In the following chapter, I will consider the so-called perceptual 'constancies', mechanisms that I argue attempt to solve stacking problems. There and in chapter 5, I will argue that 1) energy map modulations cannot solve stacking problems; 2) worldmaker contents can only provide partial solutions to stacking problems; 3) phenomenal experience cannot give determinate values for the 'stacked' properties—surface color and illumination, size and distance, e.g.; and 4) therefore, stacked properties cannot figure in phenomenal content as values, because phenomenal contents must be determinate.

Chapter 4

Constancy

In the previous chapter, I introduced a particularly difficult problem of sensory underdetermination—the problem of co-local conflation or 'stacking'. We might expect that stacked properties will appear perfectly conflated in experience, and sometimes, indeed, they do. But at other times, and maybe most of the time, we are at least partially relieved of this difficulty, and the perceptual system appears to go some of the way to dividing the stack into its constituent parts. When the system successfully divides the stack—to whatever degree—we call this 'perceptual constancy'.[1] But it turns out that it's very difficult to say precisely what is going on in such cases, certainly with respect to representational content, and even in the experience itself. In what follows, I hope to bring some clarity to the phenomenon in general by examining in detail a particular instance: the case of color perception. I choose color in part

[1] I think that 'constancy' is misleading as a description of the phenomenon, because it necessarily privileges one constituent of the stack as being the important one. We say that color constancy is constant surface representation across changes in illumination, but we could equally say that illumination is constantly represented, despite shining on different colored objects. The point is clearer in the case of size and distance—which one is held constant? Neither, really, or it depends on which is changing. But they could both be changing, and if the perceptual system does it right, then it's important only that we've pulled them apart, not which one is 'constant', since neither might be.

because it has attracted lots of attention, but also because it is in one sense the 'cleanest' of the stacking problems, because we perceive color only through a single sensory modality, whereas shape and motion are both of them perceived in multiple modalities. I do mean the arguments here to apply at least to the case of shape and size, but I will not discuss the matter here.

4.1 A Phenomenal Puzzle

Consider a uniformly-colored object, viewed partially in shade. This is about as ordinary a case as we can imagine, but describing the color look of the scene proves difficult. I will start with the following description of the case:

Puzzle: There is one sense in which the adjacent parts of the visual field (or the object) look different and another sense in which they look the same.[2]

I have tried consciously to make this description as innocent as possible, but already there is likely controversy. A staunch advocate of phenomenal 'transparency' might think it wrong to say of some part of the visual field that it 'looks' a certain way—she will say instead that only objects can look a certain way, and when they do, they look to be a certain way (hence, the parenthetical alternative). Given the rather protean nature of the 'transparency' intuition, I will ignore this objection and leave it open whether it is parts of the object or parts of the visual field that look the same (or different).[3] There is also a question about the status of 'looks', generally. I

[2] [Hilbert, 2012, p. 203] and especially [Cohen, 2010, p. 97] consider the same puzzle.

[3] One could also have the view that, e.g., parts of the visual field look different even though the objects themselves look the same (or vice-versa). We might see this as licensing a distinction between

will suppose that difference in look is phenomenal difference, and I will use the terms 'look', 'phenomenology' and 'phenomenal character' to mean the same thing.[4] But someone could contend that although there is a sameness in look, e.g., there is no sameness in phenomenal look. I cannot make sense of this distinction, and I will ignore it.

It could also be objected that one or other of the conjuncts above is false. Some will deny that there is any sense in which the parts of the visual field look the same, and one could even deny that there is any sense in which they look different. The latter seems less plausible, but not impossible, especially if one takes a very extreme stance on transparency.[5] The former actually has considerable historical precedent—it has long been thought of color constancy that it is simply a difference in look and sameness in judgment. This implies that there is no sense in which the object looks (or parts of the visual field look) uniformly-colored, despite being represented as being uniformly colored, by some faculty or other.[6] In chapter 5, I will consider a theory like this. But intuitively, there is both phenomenal sameness and difference—call these the

'objective' and 'subjective' phenomenology. Cf. [Gilchrist, 2012].

[4]A more difficult question is how to distinguish between perceptual phenomenology and other phenomenology, assuming there is any. For, if there is a distinction, we cannot tell a phenomenal difference in the color experience simply by introspecting whether there is any phenomenal difference—we'd have to know whether that difference counted as part of the color experience. This is relevant in this case, because even if we have the view that color constancy is a judgment, that judgment could itself have a phenomenal character that distinguishes it from cases without the judgment.

[5]There is a class of models for color constancy—the 'illuminant-discounting' models—that might be seen as endorsing this attitude. See, e.g., [D'Zmura and Lennie, 1986]. In fact, I think there is considerable confusion on the matter. It'd be uncharitable to imagine anyone denying that shadows are phenomenally present. Illuminant-discounting views are plausibly picking out a mechanism distinct from the one under discussion here. Such a mechanism may account for, e.g., the difference between a photograph of a bright incandescent lamp and your visual experience of it. The photograph will look very red in comparison, and this is evidence that your visual system adapts by removing some of the redness in the illuminant. Mechanisms like these probably involve low-level receptoral adaptation, and do not require direct computation of surface color. As such, they are not my target here.

[6][von Helmholtz, 1962, p. 4] has a view according to which constancies are outputs of 'unconscious conclusions'. [Reeves et al., 2008] discuss a contemporary version of the view; they call it 'projective' constancy and contrast it with 'phenomenal' constancy.

'puzzle intuitions'—and this is the puzzle I mean to solve.[7] I think that any theory of constancy should at least have something to say about why we'd have the intuition.

But why is it a puzzle? We might view it as puzzling because it seems to describe a contradiction. We ought to be able to tell whether two parts of the visual field are phenomenally the same or not, right? And this especially if we are limiting our focus to color only. How can two parts of the visual field look both same and different in color? What has gone wrong with introspection if we can't make this comparison precise? In one sense, I have already revealed part of my answer to this question. The experience of perceptual constancy is experience of an only partially-resolved sensory underdetermination problem. I am officially leaving it open whether the solution involves any phenomenal resolution, but we can make the tentative assumption that the partial solution is a partial phenomenal resolution of the problem. It remains to be seen exactly what this partial solution amounts to.

In Chapter 1, I discussed the business of explaining phenomenal experience in terms of representational content, and I am calling the contents that explain 'phenomenal contents'. In Chapter 2, I argued for some constraints on phenomenal content that issue from this commitment to explanation. We are now faced with a puzzle about phenomenal experience—how it is that two parts of the visual field can look both phenomenally different and phenomenally the same—and my commitment to explanation by phenomenal content demands that I find contents to associate with these lookings-the-same and lookings-different. We need to find, that is, a difference

[7]There is considerable evidence that subjects can make distinct sorts of color discriminations in cases like these. See, e.g., [Arend and Reeves, 1986, Troost and deWeert, 1991, Arend et al., 1991, Cornelissen and Brenner, 1995, Bäuml, 1999, Reeves et al., 2008] This does not, of course, settle the matter whether the puzzle is coherent, since discrimination is possible in the absence of phenomenology (like in blindsight, perhaps), but it is not unreasonable to suppose that the discriminative capacity is grounded in phenomenology.

of content and a sameness of content. I confront the puzzle by asking what are the relevant contents in these scenarios.

In the next section, I will describe one possible answer to the question what contents explain the phenomenal puzzle. It is the commonsense answer, and the one implied by many computational theories of color constancy currently on offer. In 4.3, I will give several arguments why it must be wrong. In 4.4, I will consider a rejoinder that involves the indeterminate representation of surface color in experience. I have argued that all phenomenal content must be determinate, and I will revisit those arguments.

4.2 The Standard View: Full Phenomenal Resolution

I said in the previous section that the pressing question for me concerns the contents that explain the phenomenal puzzle. Recall that we need both a sameness and a difference of content. In case there is a uniformly colored object partially in shade, what I will call the 'Standard View' (hereafter, 'SV') says that the sameness in content is sameness in surface reflectance, and the difference in content is difference in illumination. The original problem was sensory conflation between these two properties, and SV's solution is full phenomenal resolution. There is, of course, an easy plausibility to this view: if we know the scene to contain a uniformly-colored object and an illumination difference, and our job is to find contents to explain a phenomenal sameness and difference, these are surely the obvious candidates. Why does the part of the object in shade look different from the part in light? Answer:

that's just how uniformly-colored objects *look* in shadow. Well, not so fast. Let's back up.

SV says that surface color is a phenomenal perceptual content. Color, understood in an unanalyzed sense, is universally counted among the contents of perception, but the identification of color with surface color, though currently popular, is not universal. 'Surface color', as I am using the term here, refers to an illumination-invariant property of surfaces that is responsible (in part) for their looking the (colored) way that they do in our experience. A plausible candidate for surface color is *spectral surface reflectance*—the specific disposition of a surface to reflect light—and a popular view in both perceptual psychology and philosophy of perception is that we perceive spectral surface reflectance.

The contemporary color objectivists in philosophy hold this view [Hilbert, 1987, Matthen, 1988, Tye, 1995, Tye, 2000, Jackson, 1998, Byrne and Hilbert, 2003], and although the long history of philosophical discourse on the matter has not favored objectivism about color, recent forceful defenses of the view, as well as the current widespread popularity of intentionalism, have made color objectivism—and 'reflectance' objectivism in particular—the philosophical orthodoxy.[8]

In color science, the computational literature *defines* color constancy as the stable perception of spectral reflectance. For example, Hurlbert on the problem of color constancy: "To recover the invariant spectral reflectance properties of object surfaces from the image irradiance, in which reflectance is entangled with surface illumination." [Hurlbert, 1998, p. 283] This bias is likewise reflected in the color constancy index, a

[8]Color objectivists needn't handle the puzzle case in exactly the way I have described, but I read [Hilbert, 2005, p. 150-1] and [Tye, 2006, p. 172] as adopting that strategy. Intentionalists generally can be expected to adopt the present strategy, because they are committed to explaining phenomenology in terms of representational contents, and they typically hold that the relevant contents involve worldly, physical properties (like surface color and illumination).

measure introduced by [Arend et al., 1991, p. 665] that quantifies success in matching surfaces of the same spectral reflectance across changes in illumination.

Computational approaches to color constancy needn't but nearly universally have adopted Hurlbert's description of the problem to be solved. The vast majority propose the same solution-structure: estimate the spectral composition of the illuminant and use the estimate to compute surface color.[9] Many computational models of constancy differ only on the matter of obtaining a value for the illuminant.[10] All of them share the assumption that the system solves for reflectance and illumination, in various ways and with varying degrees of accuracy (and, perhaps, justification). Importantly for our purposes is that all such theories assume that the system computes values for these properties, whether or not the values are accurate. And it is no far stretch from this starting point to the thought that in visual experience, we see the results of these computations. And since it is not altogether implausible that the visual system might do something like this—models of the visual system can certainly do it, e.g.—SV inherits some plausibility from this corner of the empirical literature. I myself have admitted that the perceptual system generally solves underdetermination problems by what, in Chapter 3, I called 'overreach'. And this is a very ordinary such case in terms of its computational structure. The system begins with some

[9]It's worth noting that although I lump them together, [Hilbert, 2005, p. 149] finds computational approaches wanting, because he interprets them as 'illuminant-discounting' views. Recent studies have enriched this picture to include lighting dimensions among the aspects of color appearance [Logvinenko and Maloney, 2006, Tokunaga and Logvinenko, 2010, Logvinenko and Tokunaga, 2011b]. As I said in n. 6, however, I do not see 'illuminant-discounting' views as describing the phenomenon at issue here.

[10][Maloney and Wandell, 1986] are seen as the pioneers of the 'illuminant-estimating' computational approach (see also [Brainard et al., 1997, Brainard and Maloney, 2011]). Some notable examples of methods for obtaining illuminant estimates: designate a reference surface as white [Land, 1959] or grey [Land and McCann, 1971, Buchsbaum, 1980] and solve for illumination incident to those surfaces; others estimate the illuminant from specular highlights [D'Zmura and Lennie, 1986] or mutual reflections [Funt et al., 1991]. [Brainard et al., 2006] propose Bayesian methods of estimation.

assumptions, probably implicit (since perceptual), about the visible world, and it uses these to give verdicts on sensory stimuli that are ambiguous on some variables of interest. There is nothing at all wrong or implausible about this general picture of visual computation, especially as it concerns color, because our visual system has receptors that are sensitive specifically to light wavelengths, and so could plausibly perform computations on wavelength values.

In the language I've been using thus far, the standard view assumes that reflectance—or, what is easier to say, 'surface color'—and illumination figure in phenomenal content as values, and this is made possible by computations of the above sort. Furthermore, surface color and illumination figure in phenomenal content as *distinct* values, because they have been pulled apart in the computations. The 'stack' has been divided, as it were. And this means that values for surface color and illumination have each of them their own phenomenal signature—there is a certain way that every red_{37} surface looks, regardless of context, and there is a way that every $daylight_{12}$ illuminant looks, regardless of which surfaces it illuminates. If surface color and illumination are *not* conflated in experience—by way of computational solution, say—then this must be true: surface colors look a certain way, and illuminants look a certain way, and it is possible to tell them apart in experience. This explains why, in the case of the uniformly-colored object partly in shade, we notice a sameness: we notice that the surface color looks the same both in shade and in light. Likewise, it explains why we notice a difference: we notice that the object is differently illuminated across its surface. In noticing these things, we are presented with a phenomenal content according to which this part and that part of the object have red_{37} , and this part of the object is lit by $daylight_{12}$, while this one is lit by $shadow_5$.[11] In the

[11]In using these made-up descriptions of surface colors and illuminations, I am not supposing

following section, I will argue why, despite its charms, SV is wrong.

4.3 Against SV

I will offer three arguments against SV. The first establishes that surface color and illumination are conflated at least in the energy map; the second is a problem about phenomenal boundedness (and so, a problem about explanation, by my arguments); the third is a problem about behavior. The arguments are based on related considerations, and together will imply that neither surface color nor illumination can figure in phenomenal content as values (or anything like values), and that if they figure in any way, it is in the worldmaker contents.

4.3.1 The Phenomenal Trade-Off

In this section, I argue that surface color and illumination are phenomenally conflated in the energy map. SV says that surface color and illumination enjoy independent phenomenal representation. But this cannot be true, at least in the energy map. Remember that the energy map is defined by values for lightness, hue, and saturation at each location in the visual field. Now, imagine a change either to surface color or to illumination—whatever else happens in experience, surely there is a change to the energy map, and thus a change to lightness, hue, saturation, or all three, no matter whether the change is to surface or illumination.[12] Both illumination

either that the experience presents these properties under descriptions like these, or that human subjects can token surface and illumination representations at any particular level of specificity; I mean only to illustrate the fact that SV implies a discriminative capacity that might very well be described in this way.

[12]It may be possible for part of the energy map to 'return' to a previous state after chromatic or lightness adaptation, and so in particular cases this general statement might not hold through time, but ordinarily in cases of constancy, both illumination and surface changes produce changes in the

changes and surface color changes will modulate the energy map in the same way, and this is precisely because the two properties are phenomenally 'stacked'—they respond, in the same way, to the same energy changes in the sensors, and they stake out exactly the same positions in the visual field. We can say, then, that there is a kind of phenomenal 'trade-off' between surface color and illumination. A change to one makes exactly the same change to experience as a change to the other, at least in the energy map. But let's take a closer look.

Puzzle Variations

Let's consider some variations on the puzzle case. Consider two objects of the same color in the same light. SV will correctly predict phenomenal sameness all around. Now consider differently-lit objects of the same color. SV will correctly predict both phenomenal sameness and difference (after all, this is just the original case with two objects instead of one). Now consider two differently-lit objects of different colors. SV will predict in all such cases phenomenal difference all around. However, there will be certain configurations of surface color and illumination that produce phenomenal sameness. Imagine, for example, grey paper under bright light and white paper under dim light. You can adjust the lights (or swap out the sheets of paper) in such a way that you achieve phenomenal sameness in one sense.[13] And this means that SV's solution to the puzzle case cannot work for variants like these, because it cannot predict phenomenal sameness. SV is committed to explaining phenomenal sameness in terms of sameness of either surface color or illumination, but here, neither *are* the same.

---

energy map.

[13] [Chalmers, 2006, p. 85-9], [Thompson, 2006, p. 81-2], [Shoemaker, 2006, p. 474-5], [Pautz, 2009, p. 502-4] and [Foster, 2003, p. 439] make similar claims; their aims vary.

SV's explanatory failure is wider, in fact, than the set of variants like these. Imagine starting from the present case—where grey and bright look like dim and white—and swapping out the grey sheet for white. Now we have same surface, different illumination—the original puzzle case. Is there any reason to think that after swapping out the grey sheet for white we make a wholly different *type* of phenomenal comparison, the kind that SV imagines? Imagine starting from the grey-white case and *gradually* adjusting surface color until there are two white sheets. Where in that process does the *type* of phenomenal comparison change? If it doesn't change, then SV not only mishandles the grey-white case but the original puzzle case as well. SV says that in the original case, phenomenal difference is in the illuminants, and phenomenal sameness is in the surfaces. The present bit of reasoning is supposed to show that this cannot be—whether or not phenomenal sameness is sameness of surface, the difference is not difference in illumination, because that difference will give way to phenomenal sameness once we dial in grey-bright and dim-white, and in that case, we know we're not matching illuminants.

If she is to preserve the content ascriptions of her original solution, the SVer must say that there are two sorts of phenomenal difference in the puzzle case—one explained by illumination and one explained by another color content. She could then acknowledge sameness of this latter sort in the grey-white case. This move is unattractive for three reasons: it has to modify the puzzle cases in order to solve them, it modifies them by adding a third phenomenal color dimension, and it is quite contrary to the spirit of SV, whose appeal derives largely from its simplicity. SV offers a puzzle solution with a minimum of theoretical apparatus, and its content ascriptions involve only physical properties we know to be instantiated in the scene. It's tough to

see how to add to this simple picture but by giving up these advantages.

Surface-Illuminant Conflation

Plausibly, the grey-white case reveals a phenomenal conflation between surface and illumination we might not have recognized in the original puzzle case. It is well-known that surface and illumination are sensorily conflated, because one can change the spectral composition of light hitting the retina by adjusting either. The present example suggests that this conflation persists in experience. If distinct surfaces and illuminants can 'add up' to the same sums (or better, can multiply to the same products), then we must treat them, at least in one sense, as magnitudes for the same variable, and so we cannot distinguish between their respective contributions to it. In this way, surface color and illumination turn up in experience as *products* with unknown factors. Call these 'color products'.[14] Properties that enjoy entirely independent phenomenal representation are not like this. You cannot achieve phenomenal sameness by combining square and red and circle and blue, or any pair of shapes and colors, unless both the shapes and colors are themselves the same. Because this is not true of surface color and illumination, their phenomenal representation is not everywhere independent. And this means that color products are phenomenal contents (if anything is), because they can be phenomenally compared *as* products. This fact is significant, because it recommends a particular (partial) explanation of color phenomenology, and once we have that explanation in hand, we no longer appear to have any need

[14]We suppose that the factors are surface and illuminant because it is adjustments to those features that alter its appearance. This is a decidedly objectivist reading of the case, but one might alternatively incline to primitivist or relationalist [Cohen, 2009] views of the nature of the relevant color contents. I discuss the objectivist reading only because the alternatives are incompatible with SV. It's also worth pointing out that SV is implicitly a Russellian view of color contents, and one might try to avoid the conflation problem by embracing Fregeanism [Chalmers, 2006, Thompson, 2009], but even a Fregean SV is vulnerable to the arguments in 4.3.3.

for—or perhaps any room for—the explanation SV gives. That explanation depends on as-yet-unestablished assumptions about the representation of surface and illuminant and so is less attractive than the present explanation, which does not.

Importantly, however, color products alone cannot explain color phenomenology. After all, puzzle cases, including the grey-white variation, all involve both phenomenal sameness and difference, and at most, a pair of color products can account for one or the other. And nothing I've said so far implies that the additional explanans cannot be surface color.[15] And since surface color is the same in the original puzzle case and different in the grey-white case, it fits the bill.

There is trouble ahead for the hopeful SVer, however. Though it is true that color products cannot supply a complete account of color experience, their phenomenal spread is very wide. They saturate the visual field. Moreover, this aspect of color experience is very obviously value-like: its phenomenology is absolute and everywhere repeatable, points in the visual field can vary continuously in a phenomenal space defined by it, and any two points in the visual field are always comparable with respect to it, regardless of context. It is the aspect of color experience that makes natural scenes paintable onto flat, evenly-lit canvases. It is, unfortunately, the very best candidate we have in color experience for correspondence with a valued content, and surface color, if a content, is a valued content.[16] If nothing of what remains in color experience is value-apt, surface color is not a phenomenal content.

[15]I'm also not ruling out the possibility that illumination might sometimes enjoy its own independent representation. I think that it typically does not, but it is possible in certain cases to see light without (seeing) surface. Unfortunately, even in such cases, the revelation of illuminant color seems not to help us 'subtract' (phenomenally) the illuminant from the product. See [Rutherford and Brainard, 2002, Granzier et al., 2009].

[16]For, how are we to individuate, e.g., spectral surface reflectances except by values for spectral surface reflectance? Any good 'physical' color content like this is bound to individuate samely.

Alternative Interpretations of the Grey-White Case

Admittedly, the phenomenal conflation argument depends pretty significantly on a particular interpretation of the grey-white case, and an SVer might try to reject that interpretation in defense of her view. She might say, e.g., that in the grey-white case, phenomenal sameness is achieved by misrepresentation: we *mistake* grey paper under bright light for white paper under dim light (or vice-versa). The two look the same because their individual components are samely represented. If this is right, then SV needn't add surface-illuminant color products to its dimensions of color content. Alternatively, one might say that, in grey-white cases, phenomenal sameness and difference *alternate* and do not coexist. Perhaps phenomenal sameness in such cases is very fragile and very easily gives way to phenomenal difference, assuming the perceptual conditions are favorable. This may be taken to show that the search for the phenomenal presentation of surface is not at all difficult, *contra* my claim about the phenomenal 'spread' of the color products. I think, however, that both objections miss the mark.

Consider comparing two surfaces-under-illuminants through small apertures, in such a way that none of the usual cues for color constancy are available. In such conditions, we lose whatever talent we have for comparing surfaces, and we are left with a single dimension of color phenomenology. It is possible in such cases to make grey-white-type phenomenal matches, and although adjusting the conditions may upset the match, and so change the output of the comparison, they cannot change the comparison type as long as we continue to focus on the same aspect of phenomenology. And through the apertures, it is implausible to suppose we are mistaking, e.g., grey-bright for dim-white. For one thing, it would seem altogether arbitrary to choose

which among the infinity of possible configurations we are attributing to both. It is more plausible to say that we are attributing to both some shared property, and since they share neither surface nor illuminant, it can't be either of these. But this means that, whether or not grey-white matches are fragile, the point-intensities of our color phenomenology are not phenomenal presentations of either surface or illumination. Let any grey-white match give way to phenomenal difference, what results cannot be a reassignment of this particular dimension of color phenomenology to the independent comparison of surfaces (or illuminants). If we can compare surfaces, it is not by way of this aspect of color phenomenology.

This means, then, that in the energy map, at least, we've not pulled illumination out of the color product. Its difference-making contribution in the variously-lit, uniformly-colored case is as part of the phenomenal color product. It cannot also account for that difference independently, because an independent illumination can explain only a phenomenally independent illumination, and we do not have that in the energy map. Why, suppose, then, that we could possibly have surface color? Surface color is one of two parts of a color product; if we pull one out, we should get the other, and we've shown thus far that we cannot have the one. Thus, we should not have the other.

4.3.2 Phenomenal Signatures

SV implies that, by perceptual computation, we are presented in experience with values for surface color and illumination. This means that, independently of any perceptual comparisons, we can come to know about which surface properties objects have by introspection on our color experience. And so, surfaces and illuminations

must each of them have phenomenal signatures. Earlier, we ignored the computational details about constancy because they were not there terribly important. Let's take a closer look now.

Spectral surface reflectance is a function from incident light to reflected light, and the task of a constancy algorithm (according to SV) is to discover the function. But it cannot do this unless the visual system is able to make some assumptions about the function. Imagine plotting a function in 2D using only pairs of values for x and y. You can plot all day, but unless you make assumptions about some of its properties, you'll never get the function, because a function is defined by *all* of its input-output pairs. Say, for example, that you estimate the input to be {4,15,27,16,5}, and you know the output to be {2,5,9,8,5}. You might suppose, then, that the function is {0.5x, 0.33y, 0.33z, 0.5w, u}.[17] Let's grant that the input is broadband, and so we needn't worry about inputs with more set members than we've thus far encountered. There is still the problem that there are many ways to get from the input to the output. But, if we overreach and assume (more or less accurately) about surface reflectance that it just is a set of simple multipliers between 0 and 1 (as above), then we can obtain a reasonable reflectance estimate from a single pair of values. In any case, the visual system must do something like this if it is to obtain a value for reflectance from a single scene—it must do this, at least, if that value is to be anything close to the truth.[18]

This establishes that it is possible for the visual system to assign a phenomenal signature to each reflectance, or at least, if it makes no such assignment, it is not

[17]This is a function that takes, say, a vector {x,y,z,w,u} as input, and performs an element-wise multiplication, yielding a vector output.

[18]After all, it could simply *stipulate* reflectances on the basis of nothing in particular, but that's obviously not what's happening here.

for want of a feasible computational route to reflectance. But is it plausible that it makes any such assignment? We know what a phenomenal signature for reflectance would look like in the energy map—it would be the look of the surface under every possible illuminant. A signature like this cannot possibly ground one-off comparisons between items, and certainly cannot tell us anything about a surface from a single encounter. However, we do know about some objects what they look like in various lighting conditions, because we've seen them in those conditions. And we can start to make rules-of-thumb about how surfaces will change across common illumination changes. We might suppose, e.g., that the lightness look of a surface will change when we move from bright broadband light to dim broadband light. Bright reds will turn to dark reds, bright blues to dark blues, etc. And most 'natural' illumination changes are just changes in broadband lightness. But even if we have these rules of thumb in hand, do we really have, at any one moment, a phenomenal presentation of the rules? If we understand the phenomenal presentation of the rules in terms of possible color appearances, we don't really have those appearances in the present experience. We can imagine them to some degree of accuracy, but surely that can't count as figuring in phenomenal content. To see a colored surface is not to see what it would look like in all possible illumination conditions; we see it only under its present illumination.

Note that the worry here is not that reflectances are dispositions and thus cannot be visually experienced all at once (cf. [McGinn, 1996]). Let some occurrent phenomenal character signal a disposition, there is still the problem of finding the relevant phenomenal character. The point is that, plausibly, surface color as reflectance means the following for the energy map: under illuminant i_{n}, the surface will have energy map-type phenomenal character n. Reflectance, then, is the mapping from i_{n}

to n for all n. Thus, reflectance cannot be any one or even all of those characters; it is defined as the functional relation between those characters and illumination contexts. So, even if we thought that, from a single color experience, we could come to know what a surface will look like under every possible illuminant, that wouldn't be enough for the disposition itself to count as a phenomenal content. If the disposition itself is represented phenomenally, then it must have its own phenomenal signature, one over and above the energy map-type phenomenal character it guarantees for each context, and one that is the same in every context. And maybe it does. All I've tried to show thus far is that no such signature can turn up in the energy map. By definition, the energy map-type phenomenal manifestation of the disposition is different in every context.

However, we might wish to explore the possibility that surface signatures extend across viewing conditions. One could define the phenomenal signature of a surface color as the totality of its looks-under-illuminants. One might even think that some subset of its non-occurrent looks-under-illuminants are occurrently available, and that maybe this would be enough for SV here. [Cohen, 2008] has a view along these lines (though he does not identify surface colors with the colors). I doubt that anything near the totality is occurrently available, and so the property cannot strictly phenomenally be reflectance by these criteria. Much more importantly, however, it is highly counterintuitive to suppose that anything precise, beyond the occurrent look-in-illuminant, is ever available. More on the issue of representational precision later.

We might imagine phenomenal surface color to arise as a kind of imaginative subtraction of the illuminant estimate from the energy map, leaving behind the surface

color appearance. But this proposal suffers a similar fate as the last, I think, because it involves the imagination—on two counts, in fact—and therefore cannot strictly figure in phenomenal content. And even if there were a mechanism to project imagined scenes onto phenomenal experience itself, we'd be projecting imagined energy maps onto the existing one, which could do no more than produce another map.[19] As regards contents in the energy map, there is no such thing as seeing the object itself, independent of any illumination conditions and vice-versa, even though experience might at times seem this way.[20] This is just the conclusion we made in the previous section.

It follows from this that if there is to be a phenomenal signature for surface color (or illumination, for that matter), it must figure in experience as a worldmaker content. I haven't given anything like a theory of worldmaker contents, so we can't say at this point whether such contents can be values, but a worldmaker color content is already at a disadvantage in that it cannot use any aspect of the lightness, hue, or saturation spectra to present itself as a value. Introspection on color experience yields nothing obvious that fits the bill. My inability to come up with anything like a phenomenal signature for surface color in the worldmaker content could very well be a failure of introspection on my part, but I am strongly moved by the intuition that there is no one way an object looks when we judge it to have a certain surface color. Or to put the point more precisely: there seems not to be any phenomenal marker that a surface has in virtue of having the reflectance property that it does.[21] And if

[19] [Macpherson, 2012] discusses such a mechanism.

[20] Even if the object is seen in pure white light, its appearance it still sensitive to the context. Likewise for seeing an illuminant against a pure white object.

[21] And note that this extends to all kinds of perceptual contexts that affect surface color representation: center-surround, viewing angle, priming, adaptation, etc. For surface color to be represented phenomenally as a value, it would need a phenomenal marker that followed it around everywhere, the way that, for example, 2D shape phenomenology follows 2D shape properties everywhere. No

this is true, then surface color does not figure in phenomenal content as a value. We will see later that it figures in an altogether different way, if at all.

We've established thus far that surface colors and illuminations cannot figure in energy map contents (not singly, anyway), and we've suggested that surface colors and illuminations cannot figure anywhere as values. Arguments in the following section will confirm this suggestion, and it is important to keep in mind that this conclusion, though it shows SV false, nonetheless respects our initial description of the phenomenal puzzle. That puzzle concerns phenomenal comparison; nowhere in its description is the claim that color experience in optimal conditions gives us values for surface and illumination. It remains to consider whether we are capable of making cross-context comparisons between surface colors, which capability is implied by the initial description of the puzzle. If we can, it must be on the basis of something other than values for those variables—at least, it must be on the basis of something other than phenomenally presented values for those variables.

4.3.3 Matching Behavior

In the previous section, I argued that surface color is not a phenomenal content. In this section, I argue that surface color is not a perceptual content of any kind. In fact, I take my arguments here to rule out any theory according to which surface color figures in any kind of visual computation—even if it never figures in perceptual content—provided that it underwrites our surface matching ability (such as it is). I suppose it possible that the visual system computes surface color for some purpose and then ignores it when the time comes to compare surfaces, but that seems implausible.

matter the context, a 2D square will look a certain way. If 2D squarehood is instantiated, you can expect that a distinctive phenomenology is on its way. Not so for surface color, I think.

Because of this, I see my arguments here as demonstrating that the visual system never computes surface color. My own sense is that if surface color is computed, it is a content, but I am not interested in debating what counts as content.[22] I see it as an advantage of my approach here that it needn't do so.

Asymmetric Matching Accuracy

Asymmetric color matching is the dominant experimental paradigm in constancy research. An asymmetric match is one made between items under different illuminants. The subject is given control of the appearance of one of surfaces by using a knob, and is asked to make a match between this 'test' surface and a 'standard' surface whose appearance is fixed. This paradigm has been in wide use since Arend and Reeves' landmark study [Arend and Reeves, 1986]. Subject performance is measured using the color constancy index (mentioned above), which quantifies the degree of error between subject matches and perfect reflectance matches. As I said before, the nature of this measure makes clear that the presumed goal of a color constant visual system is to track surface reflectance across illumination contexts.

SV implies that asymmetric matches will be easy to make, if not also highly accurate. According to SV, we should need for matching accuracy only optimal conditions and a consistent algorithm. Even if the system misestimates illumination in an absolute sense, if it repeats the error across the divide, then the resulting comparison between surfaces should be unaffected.[23] But even for matching tasks,

[22]The computational approach is very naturally suited for and may even encourage or presuppose content ascriptions like these, in a way that other approaches in psychology are not and do not (e.g., behaviorism). See [Matthen, 1988] for discussion.

[23]In discussing their 'equivalent illuminant' model, [Brainard et al., 1997, p. 2101] admit that the illuminant estimate might be accurate at one location and inaccurate at another, within the same scene, at the same time. Thus, the error (or non-error) may not always repeat across the divide.

subject accuracy rarely exceeds 80% [Foster, 2003, p. 441]. In a highly-controlled laboratory setup, it's difficult to see how the relevant information should be available only some of the time. In natural settings, there are, of course, sharp differences in availability; in Arend and Reeves' setup (e.g.), all that changes from trial to trial is the composition of the (always broadband) illuminant and the color of the surface. We should expect that, since surface color is the supposed output of the algorithm, differences in surface color shouldn't make a difference to its operation. Surface shape and texture may make a difference, but these do not vary across trials in this study (all surfaces are flat and matte—a 'Mondrian world'). And if performance depends on the particular composition of the illuminant (if, i.e., some are easier to estimate than others), then we should see a correlation between illuminant composition and performance, but we don't. There is non-systematic variation in subject performance, variation that cuts across illumination and surface conditions.

Asymmetric Matching Certainty

This, by itself, should cast some doubt on, or at least recommend significant enrichment of, the computational models currently on offer. Subject performance is a mystery if such models correctly describe human visual computation. But rather than focus on matching accuracy, I'd like instead to focus on subject certainty. If a subject matches surfaces with certainty across illumination contexts, never realizing that her computations are in error, she should be surprised when the disappointing results are reported to her. From her perspective, the surfaces looked the same, and whatever error is explained by computational failure unavailable to her. In order for SV to escape the present threat, it must account for subject performance in this way.

For, if subjects have trouble making asymmetric comparisons, it will be difficult to maintain that their visual systems are nevertheless in possession of the information needed to make those comparisons. And if they are not, SV is false.

There is considerable evidence, in fact, that asymmetric matches are difficult or even impossible to make [Brainard et al., 1997, Logvinenko and Maloney, 2006, Tokunaga and Logvinenko, 2010, Logvinenko and Tokunaga, 2011a]. We must be careful, however, to distinguish between two claims. On the one hand, we might claim, as Brainard et al. do [p. 2098], that total phenomenal matches are impossible to make in asymmetric setups. This much seems true and is in fact built in to the puzzle cases (there is always phenomenal sameness and difference, after all) but is not to the point here. Relevant for our purposes is whether subjects can match *surfaces* across the divide, a possibility not ruled out by the former claim. [Hilbert, 2005, p. 150-1] recognizes the impossibility of total phenomenal match in asymmetric setups, but does not doubt that subjects can make asymmetric surface matches. The only caveat is that in doing so, they are matching along only a subset of the phenomenal color dimensions, of which there may be many.[24] He cannot accept, however, that this circumstance undermines the possibility of certainty in asymmetric surface matching. As an arch-objectivist, he is committed to our representing reflectances, and this means that in optimal conditions, at least, we ought to be able to match them with certainty. If it turns out that reflectance-objectivism is only tenable when modified to include conflating representations—if, i.e., the representation of reflectance is reliably

[24]Work by Tokunaga and Logvinenko [2010] suggests that color appearance includes both material and lighting dimensions. If this is true, then the objectivist can say that although asymmetric matches are impossible (since 'lighting' is among the dimensions of the comparison and is by definition unmatched), it is possible to match along material dimensions only. Unfortunately for objectivists, Tokunaga and Logvinenko themselves think that the 'illumination' dimensions nonetheless affect object appearance, and so are not strictly illumination dimensions. That is, they think that illumination and surface are partially, although not fully, phenomenally disentangled.

conflated with another kind of representation—then we don't represent reflectance. An occasionally noisy signal is one thing, but a reliably noisy signal means that we aren't getting the relevant information.[25] And if we never get it, we don't represent it.

Asymmetric Surface Matching Certainty

How, then, can we tell whether a reported impossibility (or difficulty) concerns specifically surface matches? Fortunately, a raft of studies, beginning with [Arend and Reeves, 1986], report instruction effects for asymmetric matching tasks, and their findings are quite relevant to our purpose here. In Arend and Reeves' study, subjects in one condition were asked to match the 'hue and saturation' of the 'test patch' (test surface) to the hue and saturation of a fixed standard patch.[26] In the second condition, they were asked to make the test patch 'look as if it were cut from the same piece of paper' as the standard patch [Arend and Reeves, 1986, p. 1745]. The latter condition presumably directs the subjects' attention to the surfaces, and indeed their constancy performance was better in this condition than in the first. Their methods have been repeated, sometimes with slight variations, and the findings have been similar [Arend et al., 1991, Troost and deWeert, 1991, Cornelissen and Brenner, 1995, Bäuml, 1999, Reeves et al., 2008]. So, how is this supposed to help us? If we have evidence that some matching data reflects subjects' ability to match surfaces, rather than total phenomenal appearances (as in [Brainard et al., 1997]), we have an operational distinction between surface matching

[25]This, unless the system exploits this reliability and discounts the effect of noise. But if it did, we'd experience no uncertainty.

[26]In [Arend et al., 1991, p. 664], the 'hue and saturation' match is actually a 'hue, saturation, and brightness' match. Results were similar, and I suspect the subjects were matching brightness even in the 1986 study.

and total phenomenal matching.[27]

We know, then, from the experimental results, that our talent for asymmetric surface matching is limited. The question now is whether that limitation is available to the subject. There are three distinct lines of support in favor. The first comes from [Troost and deWeert, 1991, p. 595], who report that their subjects had difficulty making matches in both conditions. This is what my arguments predict. The second line of support is several; any study that reports matching uncertainty in the asymmetric paradigm without separate instruction conditions should count as support for the claim that asymmetric surface matches are uncertain. Consider that, in the asymmetric paradigm, it is open to the subject to dial in the test settings that in fact match the standard surface. Thus, a subject given no instructions other than to make a color match can, if SV is true, simply match up the surfaces. Subjects are typically given no restrictions on what adjustments they can make to the surface; despite this fact, they report difficulty making any kind of match [Brainard et al., 1997, Logvinenko and Maloney, 2006, p. 2098]. If SV is true, why should the interposition of the other color dimensions upset their finding a surface

[27]Whether [Arend and Reeves, 1986] task was repeated by [Brainard et al., 1997], we don't know for sure. Thus, it's not clear that 'hue and saturation' matches are total phenomenal matches. Brainard et al. think so, in fact [p. 2093]. However, if, as Hilbert [Hilbert, 2005, p. 152] suggests (and [Tokunaga and Logvinenko, 2010]), there are dimensions of color appearance associated specifically with illumination, then it is not likely that a mere hue and saturation match counts as a total phenomenal match. Hilbert himself counts those particular dimensions among the 'reflectance' dimensions, and thus they cannot exhaust the entire color space. Hilbert could, of course, be mistaken (and to be fair, he's not really committed to the suggestion), and Arend and Reeves' subjects could very well have been matching on additional dimensions. [Troost and deWeert, 1991, p. 592] describe an 'exact' match condition in which subjects are asked simply to make a match 'as precisely as possible'. They identify this condition with [Arend and Reeves, 1986] 'hue and saturation' condition, and get similar results. This lends further support to the thesis that 'hue and saturation' matches are total phenomenal matches—or at least, to the thesis that their own 'exact' matches are total phenomenal matches. And interestingly, we should expect, given [Brainard et al., 1997], that Troost and de Weert's subjects will have trouble making 'exact' matches if they are total phenomenal matches—they did indeed. Other studies [Cornelissen and Brenner, 1995, Bäuml, 1999] replicated Arend and Reeves' instructions exactly, and so are no additional help here.

match?

The third and final line of support is purely introspective, so it carries less weight, but it seems appropriate here. It is worth considering that color comparison is not fraught with uncertainty in general. There are kinds of color comparison that are trivially easy. Imagine sitting at your computer, making color matches between uniformly-colored fields in Photoshop. If you make a match between two fields, and set them side by side, the result should look like an undifferentiated field of a single color, the border between the patches no longer visible. There is no phenomenal difference between the fields; they are perfectly phenomenally matched. And intuitively, it shouldn't be at all difficult to do this. Now, in this case, we could very well run up against the limits of our acuity, and there are certainly fields of slightly different color our eyes will match anyway, but the failure is not available to us. We experience no uncertainty; we don't feel like we're guessing or extrapolating or imagining. Can we say the same thing about asymmetric matches? Do subjects falter only because they are asked to make matches beyond the limits of their acuity (or because the algorithm gets it wrong)? Here is a place I think introspective evidence can be useful. Consider seeing two objects you *know* to be materially identical across a room bisected by shadow. Even in this case, knowing that they are identical, can you honestly say that they *look* that way, without question? I cannot. There is an uncertainty in the match that doesn't turn up in the Photoshop case, and that failure is available to me. Make the conditions as 'optimal' as you like, I will not trust the comparison until I bring the objects into the same illumination.

Or in what is perhaps a more natural scenario: imagine that a friend has bought some black jeans, and asks you whether they are 'the same color black' as a

sweater she's hoping to wear with them. You will instinctively bring both into the same light and make a symmetric judgment before answering. My evidence here is anecdotal, of course, but I'm confident that reflection on the matter will convince anyone that this is what they do, without even thinking about it. Now, it could be objected that the desire for a symmetric judgment is here conflated with other desires, like setting the two side-by-side, or bringing them both close to the eyes, or just getting them in good light at all, etc. And surely, illumination contexts aren't all that matters when it comes to precision, and so I admit that there is more going on here. But even if the items are side-by-side, near the face, and in bright light, if one is strongly lit by a very blue fluorescent and the other by a very yellow incandescent, you'll pull both into the glow of one or the other before you make the call. I'm sure of it.

This behavior would seem to count against SV, for obvious reasons. If that view is true, we shouldn't prefer in-context comparisons to cross-context comparisons. That is not to say that we never make cross-context comparisons. Sometimes we have to. We may, for example, wonder what the paint will look like in various lighting conditions, or how the bikini will look at the beach, despite seeing it now only under the fluorescent lights of the fitting room. In such cases, we have to do the best we can at projecting from one context to another. My point is only that this involves uncertainty, *and it shouldn't*, if SV is correct. If SV is correct, only color scientists and those in the know should have any inkling that they are regularly misrepresenting surface color across contexts. If SV is correct, surfaces wear their colors in plain sight, and even if we get the colors wrong sometimes, our judgments are certain and automatic. And this seems not to be the case.

Importantly, we have here focused on a very particular sort of asymmetric comparison, one that involves matching 'with a knob'—meaning that the end goal is a maximally precise match. Less demanding tasks can be trivially easy. Sorting objects in domains known to be restricted, like socks in a drawer, is typically easy regardless of the lighting conditions. Sorting objects by hue and saturation is often easy in natural circumstances where illumination differences are just lightness differences. Sorting objects in asymmetric conditions using coarse color categories is often easy. Shadows are easy to spot and ignore. Light can have characteristic profiles, like gradients, glare, shine, etc. These are easy to detect and ignore.[28] If these are the sorts of asymmetric comparisons we focus on, we may miss the fact that lighting conditions *do* confound our perceptions of surface. In the easy cases, representing or computing surface color is not necessary. It is only when we are asked to make color judgments that are both very certain and very precise that we even consider the reach of our discrimination capacities, and in these situations, we find them wanting. Representations of surface color are thus only necessary when we seem not to have them. What better evidence that we never represent surface color?[29]

4.4 Indeterminate Content

I anticipate the following objection here. I spent the previous section making hay of the fact that we suffer perceptual uncertainty in precisely those cases where

[28] They are easy to ignore, that is, for the purposes of the comparison. They phenomenally remain, but are ignored.

[29] There is a growing literature in color science amenable to my arguments here (most directly, [Foster, 2003]). [Amano et al., 2005] found typical asymmetric matching performance in circumstances where illuminant-estimates are not possible. This suggests that the popular computational mechanisms are otiose (see also [Granzier et al., 2009]). 'Relational' theories of color constancy work without representations or computations of surface color [Cornelissen and Brenner, 1995, Craven and Foster, 1992, Foster and Nascimento, 1994].

surface color representations would be most useful. A tidy explanation of this fact is that we nowhere represent (or even compute) surface color. An alternative explanation is that we represent surface color indeterminately—e.g., as a range of values or a graded attribution of confidence. The hypothesis correctly predicts uncertainty in the tough cases where precise surface representations are needed, and since the easy cases do not require precise surface representations, our having only imprecise surface representations is not an obstacle to perceptual certainty.

Interestingly, despite the appeal of the view, there are no well-developed theories of surface representation like this in either the philosophical or scientific literature (that I'm aware of, anyway). That asymmetric matches are difficult to make is well-known, but historically, this has persuaded commentators to 'cognitive' theories of constancy, and indeterminacy in (or suspension of) cognitive judgment is hardly worth writing about.[30] To be clear, then, at issue here are not cognitive views of constancy. Such views might make the indeterminacy proposal palatable, but I dismissed them at the outset because they very obviously deny the phenomenal puzzle around which I've structured the discussion. I think it important to respect the phenomenal intuitions that make up the puzzle, and I therefore see cognitive views as a last resort, to be considered only when we have independent reason to reject the intuitions. Moreover, it is always open to the perceiver to cognitively assess any perceptual circumstance, and in this way, cognitive theories are trivially accurate (provided they don't exclude any perceptual theories). More interesting is the question what are the final outputs

[30]Cf. [Troost and deWeert, 1991, p. 595]: "We think that since, in the object-matching condition, subjects strongly rely on indirect knowledge about objects and illuminants that have to be made explicit, the uncertainty of the matches is higher than in the exact-matching condition. Thus, in our view, the differences between the object- and exact-matching conditions are caused by different judgments rather than different sensations. If this is the case, it is simply not allowed to relate these results to color constancy as a visual phenomenon only."

of the perceptual processes. I will here consider whether those outputs might attribute surface colors indeterminately.

4.4.1 Indeterminacy and Symmetric Comparisons

I argued in Chapter 2 that phenomenal contents cannot be indeterminate. Presently, we are considering whether surface color might be represented indeterminately, and so if my arguments in Chapter 2 are successful, surface color cannot be a phenomenal content, and SV is false. I will not rehearse all of those arguments, but one in particular seems appropriate here, because I used the example of making color comparisons. There, however, I was using an un-analyzed notion of 'color', and now, we are faced with a more complicated problem, since there may very well be a difference between, in the present language, energy map comparisons and surface comparisons. If we represent color as a range or a graded attribution of confidence, the argument goes, then we cannot trust our fine-grained comparisons of color, because phenomenal sameness would indicate only that the colors both fall within the same range or that we attach the same degree of confidence to them. And that means that they could be different. But, I argue, the experience does not leave open the possibility that they could be different. Surely, the difference might lie outside our acuity—the color fields might be metamers—but even if that were true in some particular case, our experience would not tell us about it. To ascribe a range of values to an object, or to assign a degree of confidence to a representation of it, is to make a kind of epistemic confession of failure or inadequacy, and that confession, understood as *part of the very content itself*, fails to do justice to the experience.

However, it might now be objected that the argument only works when we

consider color products or Photoshop-type cases (where we compare luminous stimuli with no chance for surface-illuminant distinction). For, I myself have since argued that we don't trust our comparisons of surface color. And so, if the above argument depends on an intuition that applies only to color products, how can I use it to show that surface color contents are not represented indeterminately in experience?

In fact, it is misleading to say that we don't trust our comparisons of surface color. We do trust our *symmetric* comparisons of surface color. Color is very regularly held up as a parade case of perceptual determinacy, and for good reason: our symmetric color comparisons are maximally determinate.[31] When and only when you see two objects in the same light can you be sure whether they are maximally similar. Ascribing ranges of color values to surfaces in the same light means building into the representation the possibility that the surfaces differ in color. Since we perceive no such possibility, the indeterminacy proposal fails.[32]

To make the point more forcefully, consider that if surfaces are ascribed ranges of color values, we cannot say of any surface that it looks uniformly-colored.[33] It'd be possible, for all we'd be inclined to believe by perception alone, that the surface is a conglomeration of differently-colored points. And surely, it could be. But that possibility is not perceptually disclosed to us. So that I'm not misunderstood: I'm

[31]That is, symmetric color comparisons are maximally determinate for the representing system in which they occur. Every property is determinable relative to some standard, but that's irrelevant to the question whether the property is maximally determinate for the system in question. If two objects are perceptually identical, then whether or not they differ in fact, they are represented as being maximally similar.

[32]Note that the argument here is not threatened by the existence of visual indeterminacy in general. Assume peripheral stimuli are indeterminately represented. The fact remains that in symmetric comparisons, foveal stimuli are determinately represented. We have yet to reconcile this fact with the suggestion that in asymmetric comparisons, foveal stimuli are indeterminately represented.

[33]I've here slipped into focusing only on the ranged constructions, but only because it is syntactically smoother to consider one case instead of two. But a moment of reflection shows that, if anything, the situation is *worse* with graded attributions of confidence, because such contents are *always* attributional failures.

allowing here that surface color representations could very well *be* coarse estimations. That is, if the visual system assigns color-values to surfaces (as SV says), I have no problem with the suggestion that the value counts as coarse-grained relative to some finer-grained standard. Properties stand in whatever nomic or analytic relations to other properties that they do—the property of redness is a range of red shades. That's for sure. My point is that this sort of fact is never built in to the color experience. When two objects look the same color to you, your experience does not supplement this look with a rider that says the surfaces may in fact differ on some more precise measure of color. We *can* represent sameness of color in symmetric tasks. Where, then, could the possibility of not-sameness figure into the representation? *Ex hypothesi*, the surfaces are represented as being the same.[34]

4.4.2 Determinacy Within, Indeterminacy Across

SV now faces a dilemma: if surface color is everywhere determinate, asymmetric comparisons are inadequately explained; if surface color is everywhere indeterminate, symmetric comparisons are inadequately explained. A solution recommends itself: surface color is determinately represented within illumination contexts and indeterminately represented across them. This promises to explain both our in-context certainty and cross-context uncertainty. The following example shows this proposal deeply problematic.

Imagine that your dad has a pair of winter gloves, and you want a pair just

[34] An additional consideration in favor here: Consider that, if surfaces figured as ranges in in-context energy map comparisons, the color product should likewise be a range, but it is not presented that way. Since the color product is surface color plus illumination, we'd then have a range plus a known value (relatively speaking, at least), which would give us another range. But my initial argument about ranges in color experience rules out the color product's being a range. Thus, if you accept that argument, you have to accept this one.

like it. You go out and buy a pair you think are just like his, but alas, when you get home, you realize the pairs are not a perfect color match. One winter day, you notice your dad outside shoveling the driveway, wearing gloves. You want to determine whether he's chosen a genuine pair and not mistakenly grabbed one from each pair. You grab the two remaining gloves, and look at them side-by-side in your hand. They don't match; you're sure of it. On the present proposal, your experience ascribes determinate surface colors to the gloves, because you are in-context. Now consider holding the gloves up to the window and comparing them with the gloves your father is wearing, lit by bright sunlight. You know that his 'pair' is in fact a match for your 'pair', since you both have one of each color. And you've determined their colors inside, according to the present theory. How could it possibly be, then, that when you make the comparison through the window, the visual system no longer has available determine values for surface color? After all, you could look outside at your father's gloves, do the same computation using his illuminant (sunlight), and recover what should be the very same values. Thus, you'd have computed determinate surface colors for all four gloves. Where can indeterminacy possibly come into play here? If the first part of present proposal is correct (within-context determinacy), you ought to be able to match up the genuine pairs across the lighting context, because you've calculated the surface colors within each context. It should be as easy as matching pairs of values. But if the second part is correct, you can't make a certain match across contexts, because the surface colors are represented indeterminately when compared across contexts.

We can run the argument in the other direction, as well. Say that you compare first across the glass (an ill-advised choice, no doubt), and you're not sure about the

comparison. Why aren't you sure? Well, on the present theory, your surface color representations across the glass are indeterminate. Why? Because, presumably, your illuminant estimates are indeterminate (or noisy, or whatever), and so the computation of surface color inherits that indeterminacy. Now you come to your senses and look down at the gloves in your hand. Commonsense says that you ought to be able to tell whether the gloves in your hand are a match with one another. But we've just said that your visual computations suffer from an indeterminate illuminant estimate. If that's true, you can't even say for sure that you are within-context! You certainly can't make the comparison of surfaces in your hand. But the view promises that you can; it promises determinate in-context representation of surface. Thus, if the second part of the view is correct (cross-context indeterminacy), you can't make a certain match, and if the first part is correct, you can. The view, then, is inconsistent. What has gone wrong here?

I think the problem is a failure to respect a contextualized definition of surface color. Any theory of surface representation that says we achieve context-general surface representations must be prepared to export those representations across all contexts. Here, we've tried to do that, and finished with absurdity, because the proposal refuses to export surface representations 'across the glass', as it were. It says that within a context, the representation is determinate, while across contexts, it is indeterminate. What we ought to say if we want to respect the context-relative nature of surface color is that single contexts yield surface definitions that are not exportable across contexts. And this is why cross-context surface comparisons involve uncertainty. They mistakenly try to compare surface representations that come from incommensurable contexts.[35] My representation of the gloves inside yields comparative

[35]If 'incommensurable' sounds too strong, consider that, for precise representations of surface

data only within that context; outside the context, it is meaningless.[36] And this is true even if those representations are ranges. Consider that, within a context, we know what a range of surface colors would look like. Imagine, under a single illuminant, viewing a conveyer belt colored as a gradient. As you watch it go by, it'd be as if you were seeing the surface color change gradually—similarly under another light. But those ranges are not comparable, because if single values are incommensurate, then ranges of them are as well. Now, to be fair, my opponent here is of course proposing that surfaces across contexts are assigned ranges on the same continuum, and so they should be comparable, but we've seen already why that proposal fails: it requires that all illuminant estimates be indeterminate (because if one can be determinate, they all can) and if those are indeterminate, then no surface representation can be trusted, even within a context.[37]

---

color, commensurability *just is* illumination-invariance. If I can resolve the effect of illuminant on my surface representations in such a way as to make those representations commensurate, I've achieved illumination-invariance. I think our feeling that the contexts are comparable to some degree comes from the fact that *simpler* comparison tasks are still possible across the divide, or perhaps the fact that the demands on our perceptual imagination are often relatively light. None of this amounts to anything like illumination-invariance for spectral surface reflectance, however.

[36]This, at least with respect to the task under consideration.

[37]A lingering worry about the view just described is that it nowhere explains how we come to know that we are within a context. If it depends on making determinate illuminant estimates, why don't these estimates help us construct context-general surface representations? There are a few ways to go with this. One is to say that the system makes those estimates and solves for surface, but that these computations cannot be made phenomenally salient for one reason or another. Another possibility is that the system makes the illuminant estimate but cannot use it to compute determinate surface values, for whatever reason. Or, what seems most likely to me, is that the system doesn't make the estimates, and we assume merely that we are in-context, without assigning values to it, on the basis of other cues. I've nowhere claimed that we can't be wrong about being in-context. In fact, we surely are fooled sometimes. I've represented shadows as stains on the wall before. And this is clearly a failure of recognizing illumination context. But outside in daylight, in a clearing, with two object in your hands right in front of your face, you'll seldom be mistaken about context. And those are the parade cases of surface comparison.

4.4.3 Indeterminacy Everywhere Redux

Suppose that in response, the SVer modifies the proposal in the following way: surface color is everywhere indeterminate, and certainty in symmetric comparisons is explained by something other than surface color representation. Suppose, for example, that the visual system uses the following rule: if a comparison is symmetric (as determined by, say, illuminant cues), then total phenomenal comparison can stand in as surface comparison. I see two problems here. First, intuitively, symmetric comparisons seem like surface comparisons. At least this: they are better candidates for the role than their asymmetric counterparts. It'd be surprising if it turned out that this order is inverted. On its own, it does not seem implausible to say that total phenomenal symmetric comparisons are merely treated as precise surface comparisons. It does seem implausible to conjoin this with the claim that there are surface representations, but they are always indeterminate. For, it follows from this claim that symmetric and asymmetric comparisons are not about the same things—asymmetric comparisons are about surfaces, and symmetric comparisons are about color products (and only indirectly about surfaces). This seems implausible.

Second, the present proposal sees two distinct sorts of things going on in symmetric cases: on the one hand, there is a total phenomenal comparison that serves as a surface comparison; on the other hand, there are what must be inhibited or unattended indeterminate surface representations, waiting to be recruited for asymmetric comparisons. If they are not inhibited or unattended, then they should generate uncertainty in these cases, and we experience no such uncertainty. But the claim that they are inhibited or unattended seems implausible when we consider a compound scene that includes both symmetric and asymmetric comparisons. Imagine

three jelly beans of similar but not identical color under a lamp, two directly below and one further out, in dimmer light. You can see all three at once without shifting attention. You suffer uncertainty in the asymmetric but not the symmetric comparison. This requires that the indeterminate surface representations are simultaneously inhibited-unattended and uninhibited-attended. Impossible. A related point is that if every surface is assigned an indeterminate color, one ought to be able to attend to this content, but in symmetric comparisons, no attentional shift can generate uncertainty. The proposal fails for the same reason its predecessor failed: every asymmetric context contains symmetric contexts, and thus the structure of content attributions cannot differ between them.[38]

4.5 Final Remarks

It may be helpful to notice that the argument from surface-matching does not presuppose a division into two sorts of perceptual contents, and it establishes a far stronger claim that I need. I need only to establish that the contents that explain color experience do not involve valued representations of surface color. What the foregoing arguments conclude is that surface color (understood as valued and illumination-invariant, mind) is never represented or even computed anywhere in the perceptual system, whether it is 'phenomenal' or not. Of course, if this strong conclusion holds, the weaker does as well.

I can't decisively rule out all possible proposals involving the representation of surface color, but barring some compelling independent evidence that the visual

[38]This, at least, if we are seeking context-independent surface representation. My own considered view, as I've said, is that all surface representation is determinate but context-*dependent*.

system indeed represents surface color, the indeterminacy proposal, fraught with difficulty, seems ad hoc. After all, only on the assumption that we perceptually represent surface color does the indeterminacy proposal seem necessary. We know that perception does not give an answer to every question about surface color. But perception's not answering a question about x does not imply that it indeterminately represents x. Likewise, perception's answering a question about x does not imply that it represents x. The whole debate over the admissible contents of perception would be in vain if those principles were true. We are left, then, to choose the best explanation of our perceptual failings. In the present case, the explanatory virtues of the non-representation hypothesis outweigh the explanatory virtues of the indeterminacy hypothesis. And since SV must reject the former, SV is false.

Portions of this chapter have significant overlap with my article, published in the *Australasian Journal of Philosophy*, entitled "Surface Colour is Not a Perceptual Content" [Crockett, 2015].

Chapter 5

Denouement

In this chapter, I will describe a way the visual system might provide a 'partial' solution to the stacking problem of color without generating any indeterminate phenomenal content. This solution will be fitted to account for the pair of puzzle intuitions that a uniformly-colored object 'looks' uniformly-colored, though variously lit, and that distinct surfaces belonging to equivalent color products nonetheless 'look' different. Note that these are similar in that they signal some or other kind of separation between illuminant and surface in experience, at least to the extent that surfaces can be compared across contexts. Those particular intuitions have enjoyed very little airspace thus far, and it's time that I show how my theory can account for them without violating any of its original principles or recent conclusions. Subsequently, I will discuss the 'no-phenomenal resolution' theory—a hardline view that calls into question the intuitions themselves.

5.1 Partial Phenomenal Resolution

I have thus far concluded that surface color and illumination cannot figure independently in the energy map; they cannot figure anywhere as values; and they cannot figure anywhere indeterminately. If they figure at all independently, it can only be determinately, but not as values, in the worldmaker. This leaves little room to account for the puzzle intuitions. Before we go any further, we should consider those intuitions more closely.

5.1.1 The Puzzle Intuitions

You see a hot dog partly in shadow. It looks red all over. You see a stack of dark construction paper under your desk lamp; it doesn't look the same color as the printer paper across the desk, out of the light. These are the basic intuitions that here need explanation. Taken together, they suggest some sort of phenomenal separation between surface and illumination. There seems to be a sense in which same-colored surfaces look the same despite being differently-lit, and there seems to be a sense in which different-colored surfaces look different even when the lighting conditions are such that their color products match. The first is the parade case for color 'constancy'; the second is its inverse and seems to depend on the same phenomenal separation.

A few observations: First, we should note that although the parade case and its inverse both depend on the phenomenal separation of surface, they do so to different degrees. The parade case is supposed to involve a look of identity; its inverse involves only a look of non-identity. Thus, the parade case might seem to depend on precise computation, while its inverse need not. But I think that the parade case, as described, is risky starting point for any theory of constancy. That case, more than any other, is

at high risk of contamination from judgments that have nothing to do with the color look of the thing. If the standard view of color constancy (or some variation of it) ends up being true, then the result in the parade case should be no less certain than any asymmetric match involving samely colored objects. But we know that those cases are less certain. This suggests that neither surface sameness nor surface difference are really all that precisely rendered in experience. Our *explanandum*, then, cannot be sameness (or difference) of look across contexts. At best, it is similarity (or difference) across contexts.

And similarity will be easy to explain in some cases, because it will actually obtain in the energy map. Consider that many illumination contexts are really just lightness contexts. This is always true of sunlight, for example. And lightness contexts are easier to ignore, because they do not adjust the hue look of an object until things get very dark. So it should be no surprise that apples look red in the morning and afternoon, because the color products both in the morning and at noon are red. Not the same red, but a similar red. This is not to say that all lightness constancy can be explained on this model, but the mere fact that objects look pretty much the same across lightness contexts is not particularly impressive, and needn't depend on precise calculation of anything—it could depend on something as simple as the phenomenal similarity between color shades.

That said, we know that subjects perform above chance in surface matching tasks even when the color stimuli are very similar—when subjects 'match with the knob', so to speak—and this fact demands explanation. However, we must keep in mind that matching behavior cannot count as a direct test of the theoretical alternatives here. Subjects can and do make perceptually-grounded judgments on

the basis of more than what is directly present in the phenomenal content of the experience. So while I recognize a burden of explanation here, it is not the case that the explanation need involve, e.g., distinct phenomenal representation of surfaces. On a similar note, the intuition that we can experience phenomenal difference despite sameness in the energy map cannot depend on the observation that we can detect difference despite identical readings from a light meter, for example. The intuition can only be that we experience sameness in the energy map despite experiencing phenomenal difference. Because the energy map can involve overreach, it does not covary with readings from the light meter.[1] The moral here is that empirical test plays a secondary role to the puzzle intuitions.

We ought also to recognize a difference between the look of objects in shadow and those lit using differently-composed illuminants of similar lightness. Shadows create sharp lightness borders, whereas pairs of distinct illuminants are likely to create hue and lightness gradients, and it's certainly possible that the visual system handles these differently. Moreover, even if there were sharp illumination borders between illuminants (dividing walls, e.g.), these cases might nonetheless be distinct and differently-handled by the visual system. And, of course, it's possible that the visual system handles lightness contexts and hue contexts differently. It is often thought, for example, that extreme hue contexts—those involving narrowband illuminants, e.g.—are typically mishandled by the visual system.[2] But not so for lightness contexts.

Finally, it is important to note that the puzzle intuitions describe phenomenal comparisons. Simultaneous comparisons have been and will continue to be my target here, in large part because it'd be enormously difficult to separate out effects of

[1] See the Adelson checkerboard illusion below.

[2] Though interestingly, this is not true on my view. Since I think it is color products that explain phenomenal color experience, my view has us representing accurately more of the time, if less usefully.

perceptual memory on successive cases. We might ask about simultaneous phenomenal comparisons whether they produce phenomenal character of their own, character that we'd miss if we considered the compared objects singly. I have argued that surface color and illumination cannot figure in phenomenal content as values, which might suggest that looks of 'sameness' and 'difference' are the best we can get for surface content. Now of course, it is very rare that a visual scene offers no comparative data at all, and so perhaps reflection on ordinary experience cannot tell us much on this head. The issue is somewhat more important, I think, for the case of size and distance. The reason I mention it here is that we might be persuaded to an intuition that, for example, we represent ranges of surface colors in experience, because color experience gives us a rough idea of what surfaces will look like across contexts. However, most or all of our evidence on the matter comes from situations where comparisons are possible. And it's therefore possible that the 'range' feeling comes out of a complicated web of comparative judgments and is not a result of the way the visual system standardly computes surface values even in non-comparative cases (i.e., single-context cases).[3]

It'd be difficult to sum up everything considered here, but the distinctions will resurface in the course of giving my theory of partial phenomenal resolution. I take as a kind of primary *explanandum* the fact that we seem to experience phenomenally some degree of surface separation across contexts defined by broadband illuminants, which fact we might not expect, given what I've said about the energy map. My commitments force me to find whatever phenomenal resolution in the worldmaker, but it remains for me to establish the specific claim that phenomenal resolution is impossible in the energy map. Note that this issue is slightly different from one discussed above. There,

[3]I argued above for a kind of homogeneity between in-context and cross-context cases, but the present issue is different.

I was concerned to show that independent surface colors cannot turn up in the energy map, and the reasoning was that the energy map represents color products and so cannot also represent surface color. Now, I am concerned to show that *no* resolution of the sort we need can happen by energy map modulation, even if that resolution falls short of directly representing surface color.[4]

5.1.2 Overreach in the Energy Map

It is without question that there is overreach in the energy map. The commonest cases involve pupillary reflex and the corresponding adjustments to our experience of lightness. It is also likely that the visual system adjusts for the color temperature of broadband illuminants. You can see evidence of this by photographing indoors. An indoor scene lit with incandescents will typically look 'redder' in the photograph than through the naked eye. This suggests, then, that the visual system makes adjustments to the energy map for both lightness and hue contexts. Plausibly, then, some advantage is achieved by such adjustments. I will argue, however, that phenomenal separation of surface and illuminant cannot be achieved.

The reasoning is simple: if surface color and illumination make indistinguishable contributions to the energy map, no adjustment to the map can distinguish between the contributions. The energy map is something like a household money jar. You cannot tell, by looking at the contents of the jar, who contributed what, because contributions are made in a common currency and so are anonymous. No adjustment to the contents of the jar can change that. That said, someone with the opposite intuition might just reject the analogy.

[4]The hypothesized solution here might, for example, enable us to make accurate *judgments* about surface color despite not directly presenting surface color as a valued content.

However, it remains to consider why the visual system would make adjustments to the energy map at all if it can't do any good. The simple answer, of course, is that doing so might solve some other problem, and the adjustments have nothing to do with the phenomenal separation of surface and illumination. For example, the pupil reflex appears to be a kind of misrepresentation for the sake of more informative representation—it dims (or boosts) the lights so you can see. But not all energy map adjustments are global like this, and local adjustments likely depend on more sophisticated processes and solve more specific problems. And perhaps conflation of surface and illumination is among them. Consider the Adelson checkerboard illusion:

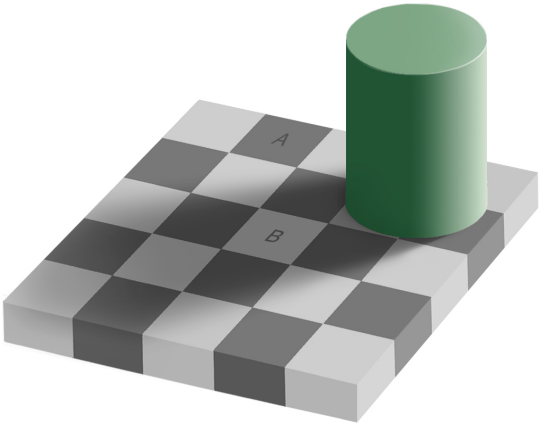

Figure 5.1: The Adelson checkerboard illusion.

Tile **A** and tile **B** are, in fact, the same color, but the visual system adjusts for lightness to achieve a certain scene interpretation (i.e., that there is a shadow and-or that **B** is a light tile). This is said to be an illusion, and my theory agrees with that analysis, but it's worth considering why it'd have to be an illusion. Whether it counts as an illusion depends, of course, on what we think the representational contents are. The ordinary interpretation of this scene is that **A** and **B** differ in surface color, and that seems to be presented in the experience. If this were a real scene, **A** and **B**

would in fact be differently colored tiles (quite differently colored, in fact), and the imposition of the shadow would darken **B** to the point where **A** and **B** give identical readings to a light meter. And this demonstration suggests that in experience, the visual system would actually boost the lightness at **B** to make sure it read as a light tile and not a dark one (or similarly, that the shadow read as a shadow). Thus, the adjustment here is to the energy map. And not only that, it seems to encourage better surface comparison. So, why is this an illusion?

I can't say for sure why the example is commonly considered an illusion, because whether it is an illusion, again, depends on your theory of content. The reason there is illusion, on my theory, is simply that the energy map represents light intensities, and so a phenomenal difference in the energy map signals a lightness difference in the world. But *ex hypothesi*, there is no lightness difference in the world, despite the fact that **A** and **B** are phenomenally different. Thus, my theory gets around the awkwardness of saying that a color illusion can promote better color representation because it distinguishes between two sorts of color representations—light intensities in the energy map and surface relations in the worldmaker. Why, then, doesn't this count as phenomenal resolution of the stacking problem of color, by way of energy map modulation?

I should say first why there is not complete phenomenal resolution of the stacking problem. Adelson's checkerboard is, of course, an asymmetric matching scenario, and so we can attempt to compare across the shadow. Compare outer tiles with their inner 'matches'. Do they appear to be exactly the same color? I think not. And this despite every cue trying to get you to believe as much (the checkerboard pattern provides a strong reason to believe there are only two tile colors).

Full phenomenal resolution could not leave the matter uncertain. Full phenomenal resolution in the energy map would be perfect covariation between the energy map and surface representation (the visual system would have to ignore illumination).[5] But this can't be what is suggested here, even though the modulation to the energy map is supposed to affect surface representation in this case (and thus, there is supposed to be at least rough covariation). For, if this were the theory, then we couldn't account for the puzzle intuitions—they report that surfaces can look to be the same (or similar) despite a difference in the energy map. We can't have it both ways—either we give up the puzzle intuitions, or we admit that we can't read off surface color from the energy map. But why doesn't the change to the energy map provide partial resolution?

To be clear, I of course concede that the energy map has something to do with surface comparisons in the scene. If the hues or saturation of tiles were wildly different, that could certainly upset the look of the surface relations amongst the tiles. So, yes, energy map modulations can act as a part of phenomenal resolution of the problem. But surface comparison cannot rest on the energy map, because the energy map is determinate everywhere, and surface comparison is not. We can compare any two points on the energy map and say whether they are phenomenally the same. But we cannot do the same with the surfaces in the scene. On my view, this is because there are no surface color representations as such in the scene, but I don't need my theory to back the claim here. Everyone should agree, regardless of theory, that surface relations

[5]There are, indeed, mechanisms to 'tune out' certain stimuli in favor of others. The auditory system tunes out droning noises so that we can better hear the novel ones, e.g.. The visual system could tune out illumination if it wanted to, in a sense, by choosing some fixed illumination level to present all surfaces (it sometimes seems to do this). Then all comparisons would be in-context. This would be like full phenomenal resolution, in a sense, except that we'd have no illumination representations at all. Clearly, however, this doesn't happen in experience, at least not most of the time. And further, it still wouldn't count as giving independent representations of both surface and illumination, and this is what I originally meant by 'full phenomenal resolution'.

are not determinate everywhere in the scene.[6] Thus, surface relations must either rely on or constitute a distinct sort of content. In my theory, this is worldmaker content. And given the little I've said about worldmaker content, it might as well be what all theories must resort to here. It is a content that is underdetermined by the energy map and so is not reducible to it. And if this is where surface comparisons reside, then the energy map cannot even share the job. For, the energy map underdetermines the worldmaker and thus can neither agree nor disagree with it directly on any matter. If it did, it would fully determine the worldmaker on that matter, and we wouldn't need the latter. The worldmaker is a content the visual system generates to get around the shortcomings of the energy map. It makes the world intelligible at a higher, more important level of description. It shares no duties with the energy map.

The preceding is a bringing together of many strains in my argument (across multiple chapters), so let me summarize in order that it be made as clear as possible. The energy map is a specification of phenomenal light intensities across the visual field. The energy map is determinate everywhere, and thus can be explained only by determinate content.[7] Relations between its contents, however, do not account for all the relations with which we seem to be presented in color experience. It is not the case that for every similarity in color experience, there is a similarity in the energy map. Likewise, it is not the case that for every dissimilarity in color experience, there

[6]Which is different, mind, than saying that surface is represented indeterminately.

[7]Not only is it a determinate content, but it is highly precise and saturates the entire visual field. There is no place in the visual field absent of an energy map content, and at every place in the visual field, it can be precisely specified. The obvious candidate for content like this is just light intensities. We know that the look of the visual field is a 'retinotopic' look—i.e., there are determinate position relations between every point on it—and thus we needn't look any further than activation across the retina for the basic picture of the energy map. There is, of course, overreach in the energy map. But it's important to see why the energy map must have contents like these—if it doesn't, then there is unexplained phenomenal character, because no other of the contents of visual experience can explain it.

is a corresponding dissimilarity in the energy map. Thus, there must be a type of color content that is underdetermined by the energy map. We have the intuition that we are presented with some surface relations in experience, and we have the intuition that these relations are underdetermined by the energy map (this is the complete set of puzzle intuitions, in fact). Thus, surface relations are plausibly this second sort of color content.

Surface relations seem, however, to be *partly* determined by the energy map. Changes in surface color manifest in the energy map. The problem is that the very same thing is true of illumination changes, and both changes manifest in precisely the same way. This means that neither can covary with the energy map, because a change to the energy map might very well mean no change to surface color, or it might mean, conversely, no change to illumination. Thus, energy map contents can be neither surface nor illumination contents, but must be a color product of both.[8] This means that surface color and illumination affect color experience both in the way they figure in the energy map, and in the way they figure in the second sort of color content, underdetermined by the energy map. I call this the 'worldmaker'. Changes to the energy map can contribute to (or what's probably more accurate, can be caused by changes in) worldmaker contents, since, for example, if you swap out an object for a darker one, the energy map will be different, too, but no change to the energy map on

[8]It can also be modulated top-down and in other ways, too, so in truth it is not fully determined by surface color and illumination either. Moreover, energy map contents cannot be presented as surfaces and illuminations, because we do not represent these determinately in experience, but we do represent energy map contents determinately. Thus, you cannot build the energy map out of independent surfaces and illuminants. It can only be accounted for by surface-plus-illumination color products. I direct the point to anyone who had thought we could account for the richness, precision, and saturation of the energy map by saying that the visual system represents surfaces and illuminants at every point. To begin, that's probably not even true, but even if it were, it could only be true in the sense that surfaces and illuminations are trapped together in the color product. Experience does not tell you 'surface A plus illuminant B'. It just gives you the sum without the addends (or, as I've been saying, the product without the factors).

its own can determine contents at this level, because again, relations in the energy map are not mappable one-one to surface relations. If the visual system adjusts the energy map despite no change in either surface color or illumination, then it misrepresents at that level (i.e., the energy map level). If the Adelson checkerboard phenomenon captures something that happens in vision, then it counts as an illusion, but it may be caused by a mechanism that works to achieve better surface representation.

5.1.3 Worldmaker Resolution

We've seen that energy map adjustments cannot determine phenomenal resolutions to stacking problems, though they may either contribute (by coaxing a certain worldmaker interpretation) or, what may be the same, participate (by being themselves adjusted by computations that determine the worldmaker). Thus, if we are to find any phenomenal resolution, we must look for it in the worldmaker.

Recall that the worldmaker is comprised of contents underdetermined by the energy map. This means that worldmaker contents cannot themselves occupy the energy map. There is no worldmaker hue, lightness, saturation, etc., because these would not be underdetermined by the energy map.[9] There are, of course, visual computations that affect these phenomenal properties, and perhaps they are common causes of both energy map and worldmaker adjustments. But contents in the worldmaker cannot themselves intrude upon the phenomenal space of the energy map. This leaves these contents little room to affect experience in striking visual ways, and their effects are more often 'felt', in a manner of speaking, than visually

[9]If this seems not inevitable, remember that the energy map is phenomenally defined. It is not defined by activation across the retina. Thus, any higher-level computational adjustment to the energy map just gives a new energy map or, we might say, is part of defining the energy map.

present. However, I take such contents still to be part of the visual look of the scene, and though they are crowded out of the primary real estate of the visual field, they are easily the most important visual contents in the system. They take our visual experience from flat fields of light points to a three-dimensional world populated by objects and surfaces, light sources, shadows and textures. Exactly how this is done cannot be my concern here; it is a job for the brain sciences. What I'd like to show is simply that, given the nature of worldmaker contents generally, full phenomenal resolution is impossible, but partial phenomenal resolution is possible by way of phenomenally-salient comparative verdicts.

Comparison in the Worldmaker

The boundaries of worldmaker content are marked on one end by the fact that they must figure in phenomenal content. This means that they must be determinate and have regular phenomenal signatures. This rules out, I think, a host of 'high-level' contents we might have thought could figure in phenomenal content because there is some unanalyzed sense in which experience can 'look to present' it that way. For example, someone can 'look like' my brother, and so I might think that the property of being my brother can be part of phenomenal content. But I think this is wrong.[10] What we mean when we say something like this is either that there are visual similarities between my brother and someone I'm seeing now, or that I am getting some visual evidence that I'm looking at my brother. Plausibly, the property of actually being my brother goes beyond what vision can present. Thus, a visually identical object needn't be nor be represented as my brother, but *ex hypothesi* could not differ with

[10]Those who think that experience presents 'singular' contents—or what is a more informative way to put it, 'object-involving' contents—may disagree here. They will allow, at least, that perception can present particulars. See, e.g., [Tye, 2011]

respect to any visual content. This shows, *inter alia*, that our unanalyzed sense that something is being 'presented' to us in experience is no guide to phenomenal content. It is possible for experience to speak indirectly and indeterminately about contents it does not itself present. This is possible because there are metaphysical relations between phenomenally presented properties and other properties. Thus, experience can seem to tell us about one by telling us about the other. Given this potential for confusion, how do we tell genuine phenomenal contents from the merely associated ones? Although there is no general answer here, I've tried to place some constraints on admissible phenomenal contents.

On the other boundary, we know that worldmaker contents cannot be so low-level that they'd feature in the energy map. And I think that this fact makes it impossible for worldmaker contents to be values for continuous variables like surface reflectance or illumination. Our ordinary notion of a 'value' for some variable likely comes from measurement using artificial scales. At least this: you probably think immediately of number when you hear 'value for a variable'. But of course, visual experience doesn't actually give us numbers for anything. There is, however, one obvious way that variables can take values in phenomenal experience: they can figure in the energy map. Fortunately for us, regarding the energy map, especially as concerns color experience, is that it is maximally phenomenally precise and determinate and so can ground precise, determinate phenomenal comparison. And perhaps more importantly, direct comparison of this sort is possible between any two points in the energy map because all points are 'valued in a common currency', phenomenally speaking. Thus, it is possible to 'measure' those contents captured in the energy map as if they were numerical values. The properties that define the dimensions of the

phenomenal space in the energy map—hue, lightness, and saturation, and location in the visual field—are given values in a fairly direct way, but it's also possible even to assign values to other properties. This second sort of assignment does not yield portable values for the variables of interest, but it does yield in-context values that are as precise and determinate as any phenomenal comparison can be (see 4.4 for the full discussion of the matter). Thus, we get a kind of 'value' for surface color (though it is an in-context only value) by comparing surfaces in single illumination contexts, and we get a 'value' for, say, height by considering only those objects against the far wall (e.g.).

This will mean that even for surface color and size or depth—properties I count out of phenomenal content, at least as values—actually have a kind of limited existence as values in phenomenal experience. And maybe this sometimes fools us into thinking that they can figure as values across contexts. But we know that these properties are not represented independently in the energy map, and thus if they are to turn up anywhere independently—and they'd have to if it's context-independence we're after—it must be in the worldmaker. But thus far, the only obvious way for any property to figure in phenomenal content as a value is through the clever manipulation of the energy map. I see no way phenomenally for the worldmaker on its own to present its contents as values for continuous variables. And this means that full phenomenal resolution is impossible in the worldmaker (where 'full resolution' is understood as division of the stack, yielding a value). And since we've shown already that it is impossible in the energy map, it is not possible anywhere. What does seem possible at this level of representation (i.e., in the worldmaker) is the phenomenal presentation of a special kind of determinate comparison, which presentation can, I argue, give a

partial phenomenal resolution of the stacking problem of color.

I began my discussion of the worldmaker in Chapter 3 with two examples of worldmaker contents—grouping and depth. Grouping is I think a perfect example of a worldmaker content. It is underdetermined by the energy map, because the energy map cannot tell us which contours belong together, it is phenomenally present, but not in the same way energy map contents are phenomenally present, and we are not at all tempted to think it exists in experience as a value of any kind. It is a 'verdict', in a sense of the word that is not linguaformal in any way (though it be may the sort of thing we could put into words), it does not involve anything like semantic and-or logical relationships, and it manifests in experience in a way that, I think, is not quite the same as the experience of having an occurrent thought or belief.[11] Grouping phenomenology is, I think, simply the phenomenal manifestation of the perceptual verdict that some contours belong to this object rather than this one (or to the same object, or whatever it is). Ambiguous stimuli help to bring out the phenomenological difference between grouping interpretations, because the interpretations are phenomenally distinct but identical in the energy map. The feeling of switching between grouping interpretations in an ambiguous scene (like in Figure 5.2 below) is, I contend, the feeling of switching between two distinct worldmaker contents. What you experience during the switch is the changing of a worldmaker content. It is difficult to view this phenomenal switch as anything other than the

[11]I do not want to take a side on the question whether there is doxastic phenomenology, but if there is such a thing, I think that worldmaker contents are different. Worldmaker contents, on my theory, explain perceptual phenomenology and so are not doxastic, and the ones I'm interested in here are supposed even further to be visual contents. To be sure, it is possible to represent surfaces in other sense modalities, and maybe there are amodal perceptual contents, but I mean the contents under discussion to be specifically visual. Worldmaker visual contents specify the ways that things can look visually over and above their energy map appearances. Two distinct worldmaker contents can have precisely the same energy map profile—this is just the observation that the energy map underdetermines the worldmaker.

manifestation of a change in perceptual verdict. When you determine (perceptually) that the arms belong to the lifter, the scene looks a certain (humorous) way. When you determine that they belong to the spotter, the scene looks a different way. Of course, the arms belong to the spotter. But it is possible to have this very sensible belief and still see the picture the other way. So, however 'top-down' the operation seems—e.g., it can be controlled voluntarily—it nonetheless seems to have nothing to do with belief. As such, I think it plausible that this is a genuine perceptual content, and in fact a phenomenal content.

Figure 5.2: Lifter. The picture is ambiguous between two grouping interpretations, one on which the hanging arms belong to the lifter and another on which they belong to the spotter behind him.

It is on this model that I propose we understand all worldmaker content. All worldmaker contents are perceptual verdicts that manifest phenomenally. And this seems reasonable, given the role that worldmaker contents play in perception. They

arise out of necessity as arbitrators on questions the energy map cannot decide on its own. Once the decision is made, a worldmaker content has no choice but to manifest phenomenally in one of two ways: (1) by modulating the energy map; or (2) by generating a visual sensation of a sort that can only be described in the terms that define the verdict. When you switch between groupings, e.g., what seems to change phenomenally is your sense of the groupings, and nothing else. Or, if there is any other sort of change, it is a change to the energy map.[12] Thus, any phenomenal change to the experience wrought by a new worldmaker interpretation is either a 'verdict-as-look' or an adjustment to the energy map, which is explained in terms of energy map contents.

Verdicts-As-Looks: The Partial Resolution

We can use the above model to understand how the visual system might make a partial phenomenal resolution of the stacking problem of color. The primary *explanandum*, again, is the intuition that we experience some degree of surface separation across illumination contexts. I have been rather cautious in my enthusiasm about our abilities here, but I want in this section to show that the general structure of the theory I'm proposing is such that it can account for a very enthusiastic interpretation of the puzzle intuitions without violating any of the central principles of my theory. The happiest interpretation of the puzzle intuitions is that in optimal conditions, we can make cross-context comparisons of surfaces as confidently as if they were in-context comparisons, and that all of the resources we need are right there in the experience.

I've argued that neither surface color nor illumination can figure in experience

[12]Tse's disks [Tse, 2005], whose appearance can be modulated voluntarily by changes in endogenous attention, are a possible example of this. Perceived brightness appears to depend on the perceived depth ordering of the disks.

as values, and to be clear, the resolution I describe here does not enable the visual system to assign values to surfaces. I've said at different points throughout what having a 'value', phenomenally, would mean. It would mean, for one, that any two surfaces assigned values could be compared phenomenally. But it would mean more than that. It would mean, further, that they can be compared using their phenomenal signatures, and this means, then, that the compared surfaces would each of them need a regular phenomenal signature. And whether or not you think that surfaces can be compared across all contexts, they do not seem to have regular phenomenal signatures. Evidently, there is need here for clarification about what it means to be comparable phenomenally.

I said two paragraphs up that the happy intuition about constancy holds that 'all of the resources we need' to make a surface comparison are 'right there in the experience'. Now, it may seem that this can mean nothing else but that the surfaces in question have regular phenomenal signatures. How else, after all, would we compare them? And I do agree that, somewhere in the visual system, there need be a comparison using values if the matter is to be settled so precisely and determinately. However, the machinery behind the comparison needn't itself belong to the experience. I am envisioning here a perceptual computation that proceeds along the lines sketched out in SV: the system estimates a value for the illuminant (a determinate value), computes surface values, and compares them. And all of this is 'sub-phenomenal', we might say. Once the verdict is made—that the two surfaces are the same or different—then the verdict and nothing else is exported into the experience in the way described in the previous section. What we experience is that the surfaces are the same or different; we do not see phenomenally whatever property it is that they share.

They look the same or different, and that is all we can say. And after all, that is all we packed into the original puzzle intuitions. On this way of phenomenal resolution, the surfaces look the same *because* they are deemed the same, not the other way around. And further, their looking the same cannot be traced to a portable 'look' that both surfaces have, that any surface would have if it were correctly represented as having whatever property the surfaces are thought to have in the background. There is no such look. All we have is the phenomenal look of sameness, and that look will turn up in any phenomenal comparison of the same sort.[13]

This might seem to do a very tidy job of accounting for the puzzle intuitions. For, we can explain phenomenal difference (or sameness) by adverting to the color product in the energy map, and we can explain phenomenal sameness (or difference) by adverting to the verdict-as-look. But there is a problem as it stands, a problem prefigured in the discussion at 4.3. It turns out that, even in optimal conditions, we cannot make precise, determinate cross-context surface matches. Those matches are not hopeless, but they are subject to error and uncertainty (and so, the error is not a matter of acuity, e.g.). This could mean many things. I have argued (at 4.4) that it cannot mean that surfaces are given in experience as indeterminate values. However, there is, manifestly, an indeterminacy to be explained. My theory rules out its being strictly a part of the experience in the sense that it'd have to be explained by an indeterminate content, but it requires some sort of explanation, and predictably, that explanation will be much like the one applied to the problem of the visual representation of my brother.

[13]We might worry then that all 'sameness' verdicts will be phenomenally identical (likewise for 'difference' verdicts). If so, then suppose that each pair of objects has additionally some unique way of looking the same. I suspect this is not necessary, since each comparison 'scene' will anyway look unique because of its constituents. The look of 'sameness' will attach to this look.

It bears reminding ourselves that the problem of indeterminacy does not affect the vast majority of cross-context comparisons. I pointed out at 4.3.3 that cross-context matches are difficult because they ask for a maximally precise judgment. They ask about identity to the highest degree the visual system can measure. But cross-context comparison need rarely worry about identity. Most of the burden of explanation here concerns the plain fact that, when I look around a variously-lit room, I am not thrown into utter hopelessness about picking apart the contributions of illuminant and surface. I will perform well above chance sorting the objects in the room according to surface color. That is an interesting fact, and it indicates that the visual system is doing something about sensory conflation. And the present view can account for it.

If we ask ourselves why asymmetric matches are difficult, it is natural to suppose that the visual system cannot be making determinate illuminant estimates. Either that or it is not applying those estimates in a determinate way to the sensor inputs. For, if it were making determinate illumination estimates (whether they were accurate or not), and applying them determinately to the sensor inputs, then we should expect at least to be able to judge of surface relations, if not see them, too. But we can't even do the first thing. It'd be a mystery, then, why the visual system would go to all the trouble of precise computation never to use the results. Similarly, it'd be odd for the system to make the determinate estimate and then fail to apply it determinately. We are left, then, with the question whether the system makes indeterminate illuminant estimates (assuming that's possible), or whether, perhaps, it makes none at all. And it may seem like the question isn't really all that important for our purposes, because whichever is true, the visual system could very well still go

about making comparisons and exporting the results into experience. After all, on the present theory, only the verdict is exported into the experience, so it really shouldn't matter how the system goes about making it. But there is one reason we might look for alternatives to the illuminant-estimating model.

In the case of color, at least, we know that the sensors measure wavelength directly (or at least, they are sensitive to specific wavelengths). And if we accept the illuminant-estimating model, then we hold that the visual system uses cues to obtain an absolute value for the illuminant (in wavelength), whether it be determinate or indeterminate. Subsequently, it uses this value to obtain a value for surface color (a set of multipliers). On the present view, this value, whether determinate or indeterminate, does not figure in experience.[14] And if it doesn't, it'd be strange for the visual system to compute it nonetheless. It's not a knock-down argument against the illuminant-estimate view, but it leaves the proponent of that view to explain the absence of the computed values in experience, an absence made especially odd by the fact that absolute values for the variables of interest could be made available to this process. She could perhaps propose that there is some basic fact about phenomenal experience at work here, but it'd be better if she didn't have to do that. It is far easier to account for the absence of these values in experience if we simply deny that the system ever computes them. That said, we have to put something in its place.

I can't run through all the computational possibilities here, but there are some general strategies that seem promising. One such strategy actually speaks directly to the puzzle intuitions. We might like to say that the hot dog partly in shadow looks

[14]The evidence for this claim comes from 4.4. A short reminder here: indeterminacy, at best, could figure in the experience as a verdict of indeterminacy. I will deny that it can, but that is irrelevant here. What we ask here is why it cannot figure in experience as an indeterminate *value*, and the answer to that, in short, is that there can be no one way that an object looks when it is attributed an indeterminate value. There can be no phenomenal signature for an indeterminate surface property.

red all over. Now, this intuition might seem to contradict some of the arguments I've made—specifically, to the claim that cross-context matches are uncertain. If the hot dog indeed looks the same all over, then it would seem that this particular cross-context match is certain. I said at 5.1.1 that the shadow case runs the risk of contamination because the parts of the surface in question belong to the same object, but perhaps there is another reason the judgment is 'contaminated', in a sense: it involves a shadow.

A shadow is not an illuminant, it can create a very abrupt illumination border, and very often its causal precursor (the obstruction) is itself visible. These facts alone might make it quite easy to detect by computation in the energy map. However shadows are detected, I think it very plausible that they are represented *as shadows* in experience. I think it very plausible, that is, that they figure in experience as worldmaker contents. The presence of a shadow, though it has a rather distinctive energy map signature, is actually underdetermined by the energy map. Shadows can look like stains and patches of spray-paint on walls, streets, grass. Try it. It is far easier to switch between shadow and 'stain' interpretations when the causal precursor of the shadow is not visible, but even then I think it possible to switch between the two interpretations. This suggests strongly that 'there is a shadow there' is a worldmaker content. If this is true, then in the hot dog case, the visual system could first judge that there is a shadow across the hot dog, see that the part in light is uniform, and judge (perhaps inaccurately) that the whole thing is one color. It could then export the verdict into the experience as a look. And it could do this all without ever making an illuminant estimate. It very naturally assigns 'shadow' to a certain energy map profile (though this assignment can be changed, as we've seen), and then it uses a

rule about the presence of shadows in making a judgment about the uniformity of the surface. I happen to think that this better explains the experience, because I cannot find the ground for the uniformity verdict in the color look of the surfaces. The best I can say about my own experience is that it looks like there is a shadow there.

We could tell a similar story about light sources. An illumination context needn't actually involve any representation of a light source. There could be a uniform glow in the room, or you could be facing away from the light source, or whatever. But we seem to be able to represent parts of a visual scene as light sources or the immediate effects of them. Glow, glare, flare, shine—these things have energy map signatures and could very well figure as worldmaker contents in experience. And even if they don't figure as contents themselves, the visual system could use their presence to identify a light source, and then apply a rule similar to shadow rule—a kind of 'discounting' of the illuminant without either representing it as a value or taking its effects out of the experience altogether. The system could, for example, represent as uniform a section of wall marked by a glare or shine.

But there are tougher cases. The two preceding involve obstructions of different sorts, but many illumination contexts are less intrusive. Taking a page from the illuminant-estimate book, we might suppose that anytime the visual system detects a color or lightness gradient in the energy map, especially one that cuts across object-boundaries (which boundaries it enforces as worldmaker content), it has identified an illuminant gradient. Objects tend not to be colored as gradients, and they certainly tend not to be arranged in such ways as to form perfect gradients across their adjacent surfaces. The visual system could, without ever computing a value for the illuminant, simply detect the illumination gradient and apply the following rules: 1) find the

object boundaries; 2) compare energy map values at the boundaries; 3) if there is difference at the boundary, the surfaces are different; 4) if there is sameness at the boundary (but a boundary is still visible); the surfaces are the same. Additionally, the system could detect whether a single object is uniformly-colored along its surface by taking several 'readings' of the energy map at adjacent points across the surface of the object, and if the differences between the points form, along some axis (say), a smooth variation in energy map color, then the object is likely to be uniformly-colored, and it thereby looks that way. Techniques like this have been proposed in the computational constancy literature, but importantly, we've described a version that involves no illumination estimate. The system simply detects that there is an illumination gradient, and it 'ignores' it, so to speak. It uses nothing more than the energy map to make its comparison. And in fact, this means that there may very well be no exportation of the verdict into the experience, since it is already manifest in the energy map.

Similarly—and this is something I discussed briefly above—the visual system might most of the time make its comparisons on the basis of actual similarity in the energy map, regardless of illumination. This is very much like the last strategy, but without even detecting anything related to the illuminant. Most illumination changes are lightness changes, and such changes, unless they involve extreme darkness, do not upset the existing relations in the energy map (other than lightness). And it is not exceedingly common to find a group of objects that differ only in lightness. Most objects in a scene are unique in terms of hue and saturation—we humans are pretty good at distinguishing these, so uniqueness is not unusual—and these relations are preserved in lightness gradients (and shadows, for that matter). Thus, the visual

system could apply these rules: 1) check for hue and saturation gradients; 2) if there are none, compare for hue and saturation, and treat these as surface color comparisons. The visual system could then either export these verdicts into the experience or not, but it likely won't need to, because the comparison is plain in the energy map. And if, in applying these rules, the visual system does find a hue gradient, it could just downregulate along the gradient the neural populations associated with these hues. Presumably, it does something like this with light bulbs, because we don't seem to 'see' all of the redness or blueness in the bulbs. This would count, of course, as an adjustment to the energy map. The system could then proceed with (2) as if the hue gradient were never there in the first place. Here, the system never makes an illuminant estimate—it finds a hue gradient, which, if part of a broadband illuminant, will make up only part of the illumination.

The little I've said cannot, of course, account for every possible situation, but in any case solutions like these seem feasible and simple, and the visual system could use combinations of them to settle a wide variety of cross-context comparisons all without making illuminant estimates.[15] It bears remembering, too, that the visual system is not thought to handle well all cross-context situations. We've been discussing strategies that work in 'optimal' conditions only, and not all conditions are optimal. And the visual system almost certainly does not have a unique strategy for every

[15]One glaring absence here is an explanation of the *same color product, different surface* cases. One problem with explaining such cases is that they'd have to be specified in more detail before my proposed strategies could be chosen among. What are the light sources like? Are they gradients? Are there shadows? Moreover, the seen difference in cases like this may actually be traceable to an energy map modulation, in which case they wouldn't count as same-color product cases. As such, I can't be sure that there truly are any such cases. I can't be sure, e.g., whether there are any cases involving two surfaces, each illuminated uniformly but distinctly, with matching color products but a difference nonetheless. And if there are clues that there is an illuminant effect, the system could use a rule like this: if there is energy map equivalence across the entire surface, but distinct illuminant effects near each surface, then they are distinct surfaces.

possible case whatever. The visual system fails sometimes, and it almost surely fails because it makes an assumption about a scene that ends up being false—it adopts a strategy that works when things are normal but not otherwise. These are the risks of overreach.

We should notice, too, that some of the foregoing strategies could very well help to ground the feeling of partial certainty (i.e., bounded uncertainty) that we get in some asymmetric matching scenarios. Part of the difficulty with the illuminant-estimate theory, as we've said, is that the indeterminate version of it can't be true, and the determinate version of it is necessarily piecemeal, because there are times when cross-context matching is uncertain (and not just inaccurate). This is not the problem of surface values in experience, but rather the problem of differentiating between 'indeterminate verdict' (not *value*, mind) and 'no verdict' cases. Sometimes, even though, *ex hypothesi*, we have computed determinate surface values, we either use them to render indeterminate verdicts about surface relations, or we don't use them at all, and we render no verdict. This is supremely weird, and it suggests that no version of the illuminant-estimate theory can be correct. And luckily, we've already considered the outlines of a replacement that doesn't suffer the same difficulties.

For one thing, nothing we've said would make it unusual for the visual system simply to withhold judgment in certain cases. If its various strategies do not render a clean result in a given case, then it leaves the matter unsettled. This could be either because none of the various strategies are actuated by the viewing conditions (if, for example, its rules are like conditionals with false antecedents), or because a given strategy yields an indeterminate result that cannot manifest in experience directly (because phenomenal contents cannot be indeterminate). Since we've not given the

visual system a unified starting point—a determinate illuminant estimate—we are not committed to seeing it through to every case.[16] This is helpful. And in the cases where we are uncertain-but-not-hopeless, we can account for our limited bit of comparative information by remembering that hue and saturation relations are often preserved across the divide. Thus, we might be unable to say whether two surfaces are identical, but we can say that both are bluish, because our visual system fails to detect any hue gradients and so makes direct cross-context comparisons of hue without regard for the details of the lightness gradient. Thus, we seem to be presented with something like a 'range of values', though for reasons I've given, this is cannot strictly be the case. What we've done instead is detect, with certainty, one aspect of similarity in the energy map profiles of the surfaces. This requires no 'ranges' and no determinate judgments of identity.

5.2 The Puzzle Intuitions Without Phenomenal Resolution

There is, however, on the view just sketched, a kind of minimization of the role of sameness and difference verdicts in experience. They play the most obvious role in the shadow or light source sorts of cases, but even in these, it might be possible to get by without the addition of the comparative verdict. Maybe representing that there is a shadow, e.g., is sufficient. In any case, I'm compelled to consider the 'no

[16] And notice that, if we think the system estimates illuminants in some cases, it must do it in every case, because on that theory, the system begins its color computations with the illuminant estimate. It cannot so much as pick out a context without it. The present proposal involving a host of strategies is similar in that if the strategy is used at all, it is used every time. However, we have defined the strategies more narrowly here, and so the conditions may fail to actuate one or many of them in particular cases. Not so for the illuminant estimate theory. There are always illumination contexts.

phenomenal resolution' (NPR) view according to which all constancy lacks phenomenal manifestation. And it anyway might be worth giving brief voice to the view for the simple reason that many people reject the puzzle intuitions as identifying a phenomenal sameness or difference in surface color. Adopting NPR does require, then, rejecting one-half of the puzzle intuitions. But I think that, perhaps, it gives a more streamlined picture of phenomenal representation and phenomenal uncertainty especially.

The difference between NPR and the view just sketched may not be all that significant, depending on how important sameness or difference verdicts-as-looks are thought to be in accounting for the puzzle intuitions. If we think the difference significant, then we don't think it possible to account for the puzzle intuitions using only energy map similarity and worldmaker representations of shadows and light sources, etc. And if we think that this makes NPR unsatisfactory as an account of color experience, then we must accept the puzzle intuitions as describing perceptual phenomenology only and not doxastic phenomenology (e.g., the experience of believing, occurrently, that two surfaces are samely-colored). It is not my intention to try to choose between the partial resolution view and NPR. I want in this section only to show how NPR might try to accommodate the puzzle intuitions without giving them a strict phenomenal reading.

First, it bears repeating a point I've made before: there is a difference between phenomenal indeterminacy and mere phenomenal underdetermination. And this is true even for variables that sometimes enjoy (determinate) phenomenal representation. Consider shape, for example. We can agree that shape—2D shape, at least—can be given determinate representation in visual experience.[17] But sometimes the shapes of

[17]Notice already that this example promises to show more than it would need to. According to my theory of surface color, surface color is never given representation in experience as such. At best, we have judgments of sameness or difference of surfaces that figure phenomenally in the experience.

objects are obscured by occluding surfaces, or smoke, or blinding light, or whatever. In these cases, the experience might not settle the question of 2D shape, though it does in optimal conditions.[18] But I think it wrong to suppose that the failure to decide 2D shape in this case means that the visual system represents shape indeterminately. We can say instead simply that it does not represent shape in this case, despite representing that there is an object there, and knowing that it must have some shape. What my theory requires is that the experience be given satisfactory explanation in terms of its phenomenal contents. And this can be achieved every bit as well using only determinate contents as it can be using indeterminate ones. I have argued, in fact, that indeterminate contents cannot explain, but I needn't repeat those arguments here. It is enough to acknowledge that we can describe the scene as one involving the following contents: an object, an occluder, the object behind the occluder, and the visible part of the object with some shape. The experience needn't say anything about the rest of the object's shape, and it really shouldn't, since it could truly be anything, so long as it fits behind the occluder. Certainly, the visual system might sometimes 'fill in the gaps' for a partially occluded object, but if it does, then this could very well count as determinate overreach. Or perhaps, if we are convinced of the 'indeterminacy' of the overreach in cases like these—and mind, any representation of an occluded part of an object involves overreach—we can account for it by appeal

Shape, on the other hand, is indeed represented in experience as such (and so, determinately). But sometimes it isn't represented at all. Thus, if it is possible in this case, we should expect it to be possible in the color case, because we know already that aspects of surface representation are regularly hidden from experience.

[18]It might even in these conditions, if we accept a phenomenal reading of object 'completion'. Two things: first, this doesn't undermine my principle concerning indeterminacy in experience (it supports that principle in fact); second, though this may seem to undermine the original premise here—that occlusion experiences don't always settle 2D shape—I think it highly implausible that object completion occurs in every case of occlusion. I think it possible to represent that there is an object behind the occluder, that it has a shape (because all objects do), but to be silent on what the shape is.

to what I will call 'imaginative overreach'.

Imaginative overreach involves the imagination's working on present perception to construct perceptual possibilities that are not strictly disclosed in the experience. As such, it is not a form of representational content in the strict sense, and does not figure in phenomenal content. I mean imaginative overreach to account for a very limited set of faint and difficult-to-discern phenomena of experience without making appeal either to perceptual contents, determinate or indeterminate, or doxastic phenomenology.[19] And as it is commonly accepted that imagination is indeed a mental faculty, that it is quasi-perceptual, and that it is distinguished from perception in part by its reduced vivacity, it shouldn't be altogether implausible that imaginative 'content' can explain such phenomena. Which phenomena exactly? It's tough to say with precision, because we might disagree on what can be accounted for by existing aspects of the view. Additionally, imagination might operate even in circumstances where whatever perceptual verdict or feeling is already accounted for in some other way. The best we can say is that imaginative overreach has the potential to help explain the sense of uncertainty in any uncertain cross-context comparison.

I do not wish to take a strong stand on theories of imagination, and so I mean for its role in the present discussion to be compatible with the major theories. For example, since my view uses imagination to explain perceptual judgment, it is neutral on the question whether imagination is pictorial.[20] For, whether or not it is pictorial, it

[19] As I've argued, I think that determinate energy map contents and certain determinate world-maker contents can account for a great deal of what we experience in cases of constancy, and that indeterminacy could very well be explained simply by the absence of determinate representation. But again, this needn't involve the presence of indeterminate representation.

[20] The contemporary notion of imagination, in the sense that I mean, is more properly called 'mental imagery', and it is an active research program in both philosophy and psychology. For a defense of the pictorial view, see [Kosslyn et al., 2006]. For a defense of the non-pictorial view, see [Pylyshyn, 2003]

is capable of delivering whatever insight explains, for example, our feeling of bounded certainty about asymmetric surface comparisons. Similarly, I think that my using imagination in the way suggested does not depend on how we understand 'Perky effects'—the finding that mental imagery can influence perception in a way analogous to masking, or, alternatively, that real stimuli can modulate mental imagery.[21] The worst-case scenario for my view is that imaginative interference with perception can generate indeterminate perceptual phenomenology, phenomenology that might very well be explained in terms of indeterminate phenomenal contents, something I explicitly deny is possible. But the Perky debate is not a debate about whether perceptual phenomenology can be indeterminate. It is a debate, most generally, about whether imaginative processes and perceptual processes spring from the same machinery, and whether they can influence each other. But even if imaginative content or phenomenology is 'overlaid' on perception, as it might be in whatever sense, perhaps depending on whether the two sorts of phenomena are compatible, it could nonetheless be the case that the resulting perceptual experience is still couched in the 'perceptual idiom', an idiom that I argue is necessarily determinate. Thus, how we interpret Perky cases needn't determine whether we think imaginative stimuli can change the basic nature of perceptual phenomenology. At the level of content, the two might influence each other, but as I have argued with respect to worldmaker contents generally, influences on perception are always translated into the perceptual idiom. This is why we see, e.g., Gestalt-switches between competing perceptual interpretations, instead of gradual easing of one interpretation into another, or through a range of them. I am here remaining silent on whether imaginative overreach might make its way in

[21]For the original study, see [Perky, 1910]. For a well-known followup study, see [Segal and Fusella, 1970]. For more recent treatments of the phenomenon, see [Craver-Lemley and Reeves, 1992] and [Reeves and Craver-Lemley, 2012].

to perceptual experience, either by generating a judgment-as-look, or by modulating perceptual phenomenology. I am committed only to saying that if it does, the resulting experience is nonetheless unchanged in its basic nature, which I argue is determinate.

Let us return to the main thread. I have argued for a way of accounting for uncertainty in cross-contexts cases on the partial resolution view. That view says that the visual system sometimes judges of surface relations and exports the results into experience. Their exportation into experience demands that they be determinate. But cross-context surface comparisons—the difficult ones, at least—often seem to involve uncertainty. If that uncertainty cannot be explained by an indeterminacy of phenomenal content, then the partial resolution view can say only that the feeling of uncertainty arises when the visual system renders no verdict. It is the feeling of a visual system accustomed to judging but withholding judgment in a particular case; or similarly, the feeling of a system accustomed to exporting its findings into experience but unable or unwilling to do so in a particular case. And this may very well work as an explanation of uncertainty here. But we might hope for a more streamlined theory of uncertainty here, one that doesn't depend on the visual system's withholding judgment (or withholding exportation) in just those cases causing problems for the view. The present suggestion, in part, is that the visual system never exports its verdicts into experience, and perhaps never renders any such verdicts directly in the first place. This makes a very natural fit, then, with my contention that no cross-context verdicts of sameness are phenomenally manifest. For, even if the judging faculty renders the matter certain—e.g., the all-over redness of the hot dog partly in shadow—that certainty is not there in the look of the thing.[22]

[22]Judgments of difference could involve certainty in the look of the thing, because the difference can be there in the energy map.

Happily, NPR is also a better fit for my theoretical strictures: we can easily account for the feeling of indeterminacy during asymmetric color tasks by adverting simply to experience-external verdicts that are themselves indeterminate. Since the verdicts in question do not figure in the phenomenal content of the experience, they needn't be determinate. If we consider again the case of 2D shape representation, we can account for the feeling of indeterminacy in the occluded case as resulting from an indeterminate verdict about shape. We are unsure about shape, and thus our judgment of shape is indeterminate. This, I think, better accounts for the feeling of indeterminacy than the 'no-verdict' account of indeterminacy in the partial resolution view.[23]

There is, however, something missing, I think, in the account as described. We have not yet explained why the faculty of perceptual judging will render judgment at all in a case of occlusion. Why doesn't it stick simply to what is phenomenally available to it? Why does it go beyond the present scene and say, 'I don't know what is behind there' or 'It might be one of these things'? On the partial resolution view, we did not have to account for the provenance of indeterminate verdicts, because on that view there are no verdicts in uncertain cases. The present account, as I said, does a better job of explaining the feeling of indeterminacy (it is explained by something rather than the absence of something), but we now have the burden of explaining why the system renders judgment at all. And here, I think, is where imaginative overreach can help explain. But let's back up a second.

The traditional form of NPR—from Helmholtz—says that sameness or similarity

[23]You might wonder why we can't just stick with the partial resolution view and attribute the indeterminacy there to an indeterminate visual judgment that doesn't make it into experience. After all, we are here describing an indeterminate judgment that likewise does not figure in experience. However, there is a difference. On NPR, there is no precedent for any judgment's figuring in experience. So, it's not figuring in here does not produce a division in the view.

and difference relations between surfaces are judged and not seen. The surfaces look just as the energy map would predict that they look, but the judging faculty is not fooled and by whatever computations or inferences is able to track constant surface properties across contexts. We know, however, that the energy map plays a more complicated role than this in perceptual constancy. For one thing, it plausibly grounds some of our cross-context comparisons of surface color. For, as we've said, surfaces typically retain aspects of similarity and difference across illumination contexts, most of which are just lightness gradients. In cases like these, no judgment is necessary to explain the look of similarity or difference. Furthermore, it is plausible that modulations to the energy map can help this very crude process along by filtering out hue differences across broadband illumination contexts (to some extent). This suggests a picture of color 'constancy' that is much less sophisticated than perhaps we had imagined, for it really involves little more than treating the energy map as a proxy for surface color and comparing 'surfaces' as if they were in the same illumination contexts. This strategy will work a surprising amount of the time, and will likewise account for the fact that we seem to get some information about how similar two surfaces are, even when we view them across-contexts. And this is precisely the machinery of the partial resolution view, minus the addition of comparative verdicts-as-looks. Thus, both the partial resolution view and NPR are ways of bringing the phenomenon of color constancy back to the energy map.[24]

Given that we've taken out of our theories of color constancy the holy world-making computational outputs of SV—i.e., values for spectral surface reflectance and illumination—and settled for something sloppy and cheap (and importantly, still phe-

[24]Although, importantly, neither view pretends to describe full phenomenal resolution in the energy map. That remains impossible.

nomenally conflated), it should be no surprise that we'd have to bring constancy back to the energy map, where surfaces and illuminations are hopelessly entangled, where lucky conditions and clever tricks are the only way to workable surface comparison. In the hard-scrabble streets of the energy map, the first rule of survival is that in-context comparisons are best. When in doubt about how to solve an asymmetric comparison, try to make it as much like an in-context comparison as possible. This rule would account for any downregulation of hue gradients in broadband illuminants, e.g. It may even account for the very notion of surface sameness as it figures in experience. After all, the only direct experience we ever have of surface sameness comes from in-context comparison. It is plausible, then, that in-context sameness is not only the gold standard for surface sameness, but the very meaning of that sameness. As such, the comparisons of any two surfaces, whether they be in-context or out, is likely guided by a desire to place them in the same context. This of course accounts for our propensity to do just that whenever we can. And when we can't, *we imagine it*.

Imaginative overreach is not like phenomenal overreach; it does not modulate phenomenal representations, and it cannot truly make the scene look any different, but it may very well affect those judgments external to experience. And I think, in fact, that it can explain them. Consider the hot dog partly in shadow. Consider the shaded part. Really look at that part. Is its color look really the same as the color look of the lit part? In any sense? (Let's ignore that it might retain some of the red hue even in shade—assume, e.g., that it is a severe shade that hides any hue information). My intuition is that it is not. However, I, like everyone, certainly represent that there is a shadow there, and it is the easiest thing in the world to judge that the entire hot dog is red. But this judgment remains a mystery if we are given nothing but the

energy map representation, plus the worldmaker representation of shadow. How do we know what shadows do to color appearance? We know it from experience, and the imagination brings it back. Picture, for example, the hot dog rolling a bit forward. Part of its surface that had been in shadow is now in light. It looks like the rest of the surface in light. It's no great leap to imagine it rolling entirely out of shadow, and looking red all over. And it is on these grounds that I think we judge the thing uniform.

A similar process could occur in the 'indeterminate' cases. Remember the black sweaters in lamplight? We will naturally want to compare their colors under the same lamp, but if we were prevented from doing do, we'd imagine it anyway. We'd see what one looks like under the incandescent and we'd imagine what the other looks like under the incandescent, even though we cannot move it from under the fluorescent.[25] We imagine moving it because in moving it we give it a phenomenal signature—a non-portable signature, of course, but as good as any we can get. And out of this imaginative process comes not a phenomenal representation of sameness or difference, but a verdict of sameness or difference, guided by the operation of the imagination. And in cases like these, the verdict must remain indeterminate. We can't be sure of the color look under a new illuminant, but we have some idea, drawn from experience and re-created to the best of our imaginative powers in the present case. The role of imaginative overreach in cases like these helps account for overreach in the verdict. Without the imagination, the judging faculty would have no reason to reach beyond the present experience and guess at what might be, were the conditions more favorable.

[25] [Cohen, 2008] gives an account of color constancy much like this, but he sees the counterfactual as specifying contents. I am not going as far as that. At least, I am not claiming that the counterfactual specifies a phenomenal content. On his view, the counterfactual is not generated by the imagination, and so it does not have that obstacle to figuring as a content.

With the imagination, guided by memories of color experience past, it's difficult to see how the judging faculty could possibly resist reaching beyond the present experience.

It is in this way that the role of imaginative overreach makes the mechanism complete. It arises out of a very plausible propensity to rely on past color experience, to place things in like illumination contexts, even if they cannot be moved, and it accounts for our propensity to judge of what is not strictly there to be seen. And although the executive powers of the system can wrest control of the imagination, it is unlikely that imaginative overreach is typically voluntary. This makes perceptual judgment seem a bit less agent-controlled or cognitive, which in cases like these is a good thing, because surely the 'judgments' in question can remain even if we are told that we are in the grip of illusion. In this way, the sort of judging at issue here is not cognitive judgment and needn't have any necessary connection to belief. Moreover, the present view is preferable to a view which accounts for the feeling of indeterminacy by appeal to doxastic phenomenology. That view suffers doubly: it depends on the existence of doxastic phenomenology, and it leaves unexplained the origin of the judgment. Like the partial resolution view, it makes the experience of uncertainty or indeterminacy (or certainty, depending on the case) a consequence of the judgment, which in itself is not objectionable, but then it forgets to say why the system judges as it does in the first place. The imaginative overreach account helps to bridge the gap between experience and judgment in a way that gets right the inferential ordering. It is for these reasons that I see NPR, supplemented as I've done with my brief account of imaginative overreach, as a powerful opponent to the partial resolution view, despite its failing to account for a strict perceptual reading of the puzzle intuitions.

5.3 Summary and Concluding Remarks

This concludes my account of partial solution to the stacking problem of color. I have argued that the most direct and initially appealing theory of phenomenal resolution must be false; surface color and illumination never figure in phenomenal content as values. Seeing that these are left out of experience, the computational machinery thought to produce them is unlikely to exist. In its place, I've proposed two distinct accounts of partial solution to the stacking problem of color. Both respect my constraints on phenomenal contents—that they be determinate and phenomenally unique. They differ on the question whether the puzzle intuitions–that there exists some representational separation of surface color and illumination—are to be given a strict reading, such that the separation is phenomenally salient.

We began the project with the idea that perception is *presentational* and not *descriptive* and that this places significant constraints on what sorts of contents can figure in perceptual representations. And we've now come to understand what this means in more specific terms. The twin constraints of determinacy and uniqueness are direct consequences of the fact that perception is presentational. If a content cannot demand phenomenal uniqueness, then it cannot be presented, and thus it cannot be a perceptual content. If a content is indeterminate, it cannot be presented, and thus it cannot be a perceptual content. And although, as we've seen, perception overreaches its sensor input, the resulting phenomenal contents must all of them be presentable. In the specific case of color, surface and illumination cannot be presented independently as values, so they cannot be perceptual contents. So, despite the fact that all of the contents my constraints rule out can be *described* and even *computed*, they are never perceptual because they cannot be shown. Precisely why this is remains a point of

some mystery. It is difficult to say, in particular, what constrains presentation in the worldmaker, but it does indeed appear constrained. In the contemporary debate over admissibility, worldmaker-apt contents are now attracting considerable attention. It is my hope that the foregoing has given reason to favor some among the positions in that debate.

Bibliography

[Amano et al., 2005] Amano, K. D., Foster, D. H., and Nascimento, S. M. C. (2005). Minimalist surface-colour matching. *Perception*, 34:1007–1011.

[Arend and Reeves, 1986] Arend, L. and Reeves, A. (1986). Simultaneous color constancy. *Journal of the Optical Society of America A*, 3(10):1743–1751.

[Arend et al., 1991] Arend, L., Reeves, A., Schirillo, J., and Goldstein, R. (1991). Simultaneous color constancy: patterns with diverse Munsell values. *Journal of the Optical Society of America A*, 8:661–672.

[Ayer, 1963] Ayer, A. J. (1963). *The Foundations of Empirical Knowledge*. MacMillan, London.

[Barlow, 1959] Barlow, H. B. (1959). Sensory mechanisms, the reduction of redundancy, and intelligence. *The mechanisation of thought processes*, 10:535–539.

[Bäuml, 1999] Bäuml, K.-H. (1999). Simultaneous colour constancy: how surface color perception varies with the illuminant. *Vision Research*, 39(8):1531–1550.

[Bayne, 2009] Bayne, T. (2009). Perception and the reach of phenomenal content. *The Philosophical Quarterly*, 59(236):385–404.

[Block, 1990] Block, N. (1990). Inverted earth. *Philosophical Perspectives*, 4:53–79.

[Block, 1996] Block, N. (1996). Mental paint and mental latex. In Villenueva, E., editor, *Philosophical Issues*, volume 7. Ridgeview Publishing Company, Atascadero, California.

[Block, 2010] Block, N. (2010). Attention and mental paint. *Philosophical Issues*, 20(1):23–63.

[Brainard et al., 1997] Brainard, D. H., Brunt, W. A., and Speigle, J. M. (1997). Color constancy in the nearly natural image I. asymmetric matches. *Journal of the Optical Society of America, A*, 14(9):2091–2110.

[Brainard et al., 2006] Brainard, D. H., Delahunt, P. B., Freeman, W. T., Kraft, J. M., and Xiao, B. (2006). Bayesian model of human color constancy. *Journal of Vision*, 6:1267–1281.

[Brainard and Maloney, 2011] Brainard, D. H. and Maloney, L. T. (2011). Surface color perception and equivalent illumination models. *Journal of Vision*, 11(5):1.

[Brewer, 2006] Brewer, B. (2006). Perception and content. *European Journal of Philosophy*, 14(2):165–181.

[Bruner et al., 1949] Bruner, J. S., Postman, L., et al. (1949). On the perception of incongruity: A paradigm. *Journal of personality*, 18(2):206–223.

[Buchsbaum, 1980] Buchsbaum, G. A. (1980). A spatial processor model for object colour perception. *Journal of the Franklin Institute*, 310:1–26.

[Byrne, 2002] Byrne, A. (2002). Intentionalism defended. *The Philosophical Review*, 110:199–240.

[Byrne, 2009] Byrne, A. (2009). Experience and content. *Philosophical Quarterly*, 59(236):429–451.

[Byrne and Hilbert, 2003] Byrne, A. and Hilbert, D. R. (2003). Color realism and color science. *Behavioral and Brain Sciences*, 26(1):3–64.

[Byrne and Tye, 2006] Byrne, A. and Tye, M. (2006). Qualia ain't in the head. *Noûs*, 40(2):241–255.

[Caclin et al., 2012] Caclin, A., Paradis, A.-L., Lamirel, C., Thirion, B., Artiges, E., Poline, J.-B., and Lorenceau, J. (2012). Perceptual alternations between unbound moving contours and bound shape motion engage a ventral/dorsal interplay. *Journal of vision*, 12(7):11.

[Campbell, 2002] Campbell, J. (2002). *Reference and Consciousness*. Oxford University Press.

[Chalmers, 2006] Chalmers, D. (2006). Perception and the fall from Eden. In Gendler, T. S. and Hawthorne, J., editors, *Perceptual Experience*, pages 49–125. Oxford University Press, New York.

[Clark, 2013] Clark, A. (2013). Whatever next? predictive brains, situated agents, and the future of cognitive science. *Behavioral and Brain Sciences*, 36(03):181–204.

[Cohen, 2008] Cohen, J. (2008). Color constancy as counterfactual. *Australasian Journal of Philosophy*, 86(1):61–92.

[Cohen, 2009] Cohen, J. (2009). *The Red and The Real: An Essay on Color Ontology*. Oxford University Press, Oxford.

[Cohen, 2010] Cohen, J. (2010). Perception and computation. *Philosophical Issues*, 20(1):96–124.

[Cornelissen and Brenner, 1995] Cornelissen, F. W. and Brenner, E. (1995). Simultaneous colour constancy revisited: an analysis of viewing strategies. *Vision Research*, 35:2431–2448.

[Craven and Foster, 1992] Craven, B. J. and Foster, D. H. (1992). An operational approach to colour constancy. *Vision Research*, 32(7):1359–1366.

[Craver-Lemley and Reeves, 1992] Craver-Lemley, C. and Reeves, A. (1992). How visual imagery interferes with vision. *Psychological review*, 99(4):633.

[Crockett, 2015] Crockett, D. (2015). Surface colour is not a perceptual content. *Australasian Journal of Philosophy*, 93(2):303–318.

[Dennett, 1991] Dennett, D. C. (1991). *Consciousness Explained*. Little, Brown and Company, Boston.

[Dretske, 1995] Dretske, F. (1995). *Naturalizing the Mind*. MIT Press, Cambridge, Massachusetts. Originally delivered as the 1994 Jean Nicod Lectures.

[Dretske, 1981] Dretske, F. I. (1981). *Knowledge and the Flow of Information*. MIT Press, Cambridge, Massachusetts.

[D'Zmura and Lennie, 1986] D'Zmura, M. and Lennie, P. (1986). Mechanisms of color constancy. *Journal of the Optical Society of America A*, 3:1662–1672.

[Farennikova, 2013] Farennikova, A. (2013). Seeing absence. *Philosophical studies*, 166(3):429–454.

[Foster, 2003] Foster, D. H. (2003). Does colour constancy exist? *Trends in Cognitive Science*, 7(10):439–443.

[Foster and Nascimento, 1994] Foster, D. H. and Nascimento, S. M. C. (1994). Relational colour constancy from invariant cone-excitation ratios. *Proceedings of the Royal Society of London B*, 257:115–121.

[Friston, 2005] Friston, K. (2005). A theory of cortical responses. *Philosophical transactions of the Royal Society B: Biological sciences*, 360(1456):815–836.

[Funt et al., 1991] Funt, B., Drew, M., and Ho, J. (1991). Color constancy from mutual reflection. *International Journal of Computer Vision*, 6:5–24.

[Gelman and Markman, 1987] Gelman, S. A. and Markman, E. M. (1987). Young children's inductions from natural kinds: The role of categories and appearances. *Child development*, pages 1532–1541.

[Gentner, 1983] Gentner, D. (1983). Structure-mapping: A theoretical framework for analogy. *Cognitive science*, 7(2):155–170.

[Gilchrist, 2012] Gilchrist, A. (2012). Objective and subjective sides of perception. *Visual Experience: Sensation, Cognition, and Constancy*, page 105.

[Granzier et al., 2009] Granzier, J. J., Brenner, E., and Smeets, J. B. (2009). Can illumination estimates provide the basis for color constancy? *Journal of Vision*, 9(3):18.

[von Helmholtz, 1962] von Helmholtz, H. (1962). *Helmholtz's Treatise on Physiological Optics*. Dover, New York. Originally published in 1867.

[Hilbert, 2012] Hilbert, D. (2012). Constancy, content, and inference. *Visual Experience: Sensation, Cognition and Constancy*, pages 199–211.

[Hilbert, 1987] Hilbert, D. R. (1987). *Color and Color Perception: A Study in Anthropocentric Realism*. CSLI, Stanford.

[Hilbert, 2005] Hilbert, D. R. (2005). Color constancy and the complexity of color. *Philosophical Topics*, 33:141–158.

[Hill, 2009] Hill, C. S. (2009). *Consciousness*. Cambridge University Press.

[Horgan and Tienson, 2002] Horgan, T. and Tienson, J. (2002). The phenomenology of intentionality and the intentionality of phenomenology. In Chalmers, D. J., editor, *Philosophy of Mind: Classical and Contemporary Readings*, pages 520–533. Oxford University Press.

[Hupe et al., 1998] Hupe, J., James, A., Payne, B., Lomber, S., Girard, P., and Bullier, J. (1998). Cortical feedback improves discrimination between figure and background by v1, v2 and v3 neurons. *Nature*, 394(6695):784–787.

[Hurlbert, 1998] Hurlbert, A. C. (1998). Computational models of color constancy. In Walsh, V. and Kulikowski, J., editors, *Perceptual Constancy: Why Things Look as They Do*, pages 283–322. Cambridge University Press, Cambridge.

[Jackson, 1982] Jackson, F. (1982). Epiphenomenal qualia. *Philosophical Quarterly*, 32(127):127–136.

[Jackson, 1998] Jackson, F. (1998). *From Metaphysics to Ethics: A Defence of Conceptual Analysis*. Oxford, New York. Originally given as the 1998 Locke Lectures.

[Jackson, 2003] Jackson, F. (2003). Mind and illusion. *Royal Institute of Philosophy Supplement*, 53:251–271.

[Kosslyn et al., 2006] Kosslyn, S. M., Thompson, W. L., and Ganis, G. (2006). *The case for mental imagery.* Oxford University Press.

[Lamme and Roelfsema, 2000] Lamme, V. A. and Roelfsema, P. R. (2000). The distinct modes of vision offered by feedforward and recurrent processing. *Trends in neurosciences*, 23(11):571–579.

[Land, 1959] Land, E. H. (1959). Color vision and the natural image. part i. *Proceedings of the National Academy of Sciences of the United States of America*, 45(1):115.

[Land and McCann, 1971] Land, E. H. and McCann, J. J. (1971). Lightness and retinex theory. *Journal of the Optical Society of America*, 61.

[Lee and Mumford, 2003] Lee, T. S. and Mumford, D. (2003). Hierarchical bayesian inference in the visual cortex. *JOSA A*, 20(7):1434–1448.

[Lerner et al., 2001] Lerner, Y., Hendler, T., Ben-Bashat, D., Harel, M., and Malach, R. (2001). A hierarchical axis of object processing stages in the human visual cortex. *Cerebral Cortex*, 11(4):287–297.

[Logvinenko and Maloney, 2006] Logvinenko, A. D. and Maloney, L. T. (2006). The proximity structure of achromatic surface colors and the impossibility of asymmetric lightness matching. *Perception & Psychophysics*, 68(1):76–83.

[Logvinenko and Tokunaga, 2011a] Logvinenko, A. D. and Tokunaga, R. (2011a). Colour constancy as measured by least dissimilar matching. *Seeing and perceiving*, 24(5):407–452.

[Logvinenko and Tokunaga, 2011b] Logvinenko, A. D. and Tokunaga, R. (2011b). Lightness constancy and illumination discounting. *Attention, Perception, & Psychophysics*, 73(6):1886–1902.

[Lycan, 1996] Lycan, W. G., editor (1996). *Consciousness and Experience*. MIT Press, Cambridge, Massachusetts.

[Macpherson, 2012] Macpherson, F. (2012). Cognitive penetration of colour experience: Rethinking the issue in light of an indirect mechanism. *Philosophy and Phenomenological Research*, 84(1):24–62.

[Maloney and Wandell, 1986] Maloney, L. T. and Wandell, B. A. (1986). Color constancy: a method for recovering surface spectral reflectance. *Journal of the Optical Society of America A*, 3(1):29–33.

[Matthen, 1988] Matthen, M. (1988). Biological functions and perceptual content. *The Journal of Philosophy*, 85:5–27.

[McGinn, 1991] McGinn, C. (1991). *The Problem of Consciousness*. Basil Blackwell, Oxford.

[McGinn, 1996] McGinn, C. (1996). Another look at color. *The Journal of Philosophy*, 93(11):537–553.

[Mevo et al., 2013] Mevo, C., Kourtzi, Z., and Tsal, Y. (2013). Expectation induced curvature perception in v2. *Journal of Vision*, 13(9):502–502.

[Moore, 1903] Moore, G. E. (1903). The refutation of idealism. *Mind*, pages 433–453.

[Morrison, 2015] Morrison, J. (2015). Perceptual confidence. *Analytic Philosophy*.

[Mumford, 1991] Mumford, D. (1991). On the computational architecture of the neocortex. *Biological cybernetics*, 65(2):135–145.

[Murray et al., 2002] Murray, S. O., Kersten, D., Olshausen, B. A., Schrater, P., and Woods, D. L. (2002). Shape perception reduces activity in human primary visual cortex. *Proceedings of the National Academy of Sciences*, 99(23):15164–15169.

[Pascual-Leone and Walsh, 2001] Pascual-Leone, A. and Walsh, V. (2001). Fast back-projections from the motion to the primary visual area necessary for visual awareness. *Science*, 292(5516):510–512.

[Pautz, 2009] Pautz, A. (2009). What are the contents of experiences? *The Philosophical Quarterly*, 59(236):483–507.

[Peacocke, 1983] Peacocke, C. (1983). *Sense and Content*. Oxford University Press, Oxford.

[Perky, 1910] Perky, C. W. (1910). An experimental study of imagination. *The American Journal of Psychology*, pages 422–452.

[Pylyshyn, 1999] Pylyshyn, Z. W. (1999). Is vision continuous with cognition? the case for cognitive impenetrability of visual perception. *Behavioral and Brain Sciences*, 22(3):341–423.

[Pylyshyn, 2003] Pylyshyn, Z. W. (2003). *Seeing and Visualizing: It's Not What You Think*. MIT Press, Cambridge, Massachusetts.

[Rao and Ballard, 1999] Rao, R. P. and Ballard, D. H. (1999). Predictive coding in the visual cortex: a functional interpretation of some extra-classical receptive-field effects. *Nature neuroscience*, 2(1):79–87.

[Reeves and Craver-Lemley, 2012] Reeves, A. and Craver-Lemley, C. (2012). Unmasking the perky effect: spatial extent of image interference on visual acuity. *Frontiers in psychology*, 3.

[Reeves et al., 2008] Reeves, A. J., Amano, K., and Foster, D. H. (2008). Color constancy: Phenomenal or projective? *Perception & psychophysics*, 70(2):219–228.

[Russell, 1912] Russell, B. (1912). *The Problems of Philosophy*. Oxford University Press, London.

[Rutherford and Brainard, 2002] Rutherford, M. D. and Brainard, D. H. (2002). Lightness constancy: a direct test of the illumination-estimation hypothesis. *Psychological Science*, 13(2):142–149.

[Schellenberg, 2011] Schellenberg, S. (2011). Perceptual content defended. *Noûs*, 45(4):714–750.

[Segal and Fusella, 1970] Segal, S. J. and Fusella, V. (1970). Influence of imaged pictures and sounds on detection of visual and auditory signals. *Journal of experimental psychology*, 83(3p1):458.

[Shoemaker, 2006] Shoemaker, S. (2006). On the way things appear. In Gendler, T. S. and Hawthorne, J., editors, *Perceptual Experience*, pages 461–480. Oxford University Press, New York.

[Siegel, 2010] Siegel, S. (2010). *The Contents of Visual Experience*. Oxford University Press, New York.

[Siewert, 1998] Siewert, C. (1998). *The significance of consciousness*. Princeton University Press.

[Stazicker, 2011] Stazicker, J. (2011). Attention, visual consciousness and indeterminacy. *Mind & Language*, 26(2):156–184.

[Stazicker, 2014] Stazicker, J. (2014). The visual presence of determinable properties. In Dorsch, F., editor, *Phenomenal Presence*. Oxford University Press.

[Stoljar, 2004] Stoljar, D. (2004). The argument from diaphanousness. *Canadian Journal of Philosophy*, 34(sup1):341–390.

[Stoner and Albright, 1992] Stoner, G. R. and Albright, T. D. (1992). Neural correlates of perceptual motion coherence. *Nature*, 358(6385):412–414.

[Thompson, 2006] Thompson, B. (2006). Colour constancy and Russellian representationalism. *Australasian Journal of Philosophy*, 84(1):75–94.

[Thompson, 2009] Thompson, B. (2009). Senses for senses. *Australasian Journal of Philosophy*, 87(1):99–117.

[Tokunaga and Logvinenko, 2010] Tokunaga, R. and Logvinenko, A. D. (2010). Material and lighting dimensions of object colour. *Vision Research*, 50(17):1740–1747.

[Travis, 2004] Travis, C. (2004). The silence of the senses. *Mind*, 113(449):57–94.

[Troost and deWeert, 1991] Troost, J. M. and deWeert, C. M. M. (1991). Naming versus matching in color constancy. *Perception & Psychophysics*, 50:591–602.

[Tse, 2005] Tse, P. U. (2005). Voluntary attention modulates the brightness of overlapping transparent surfaces. *Vision research*, 45(9):1095–1098.

[Tye, 1995] Tye, M. (1995). *Ten Problems of Consciousness: A Representational Theory of the Phenomenal Mind*. MIT Press, Cambridge, Massachusetts.

[Tye, 2000] Tye, M. (2000). *Consciousness, Color, and Content*. MIT Press, Cambridge, Massachusetts.

[Tye, 2006] Tye, M. (2006). In defense of representationalism: Reply to commentaries. *Pain: New Essays on Its Nature and the Methodology of Its Study, Cambridge, MA, The MIT Press, Bradford Books*, pages 163–176.

[Tye, 2011] Tye, M. (2011). *Consciousness revisited: Materialism without phenomenal concepts*. Mit Press.

[Williams et al., 2008] Williams, M. A., Baker, C. I., de Beeck, H. P. O., Shim, W. M., Dang, S., Triantafyllou, C., and Kanwisher, N. (2008). Feedback of visual object information to foveal retinotopic cortex. *Nature neuroscience*, 11(12):1439–1445.

[Wokke et al., 2012] Wokke, M. E., Vandenbroucke, A. R., Scholte, H. S., and Lamme, V. A. (2012). Confuse your illusion feedback to early visual cortex contributes to perceptual completion. *Psychological science*, page 0956797612449175.