

Perception and the Fall from Eden

David J. Chalmers
Philosophy Program
Research School of Social Sciences
Australian National University

1 Eden

In the Garden of Eden, we had unmediated contact with the world. We were directly acquainted with objects in the world and with their properties. Objects were presented to us without causal mediation, and properties were revealed to us in their true intrinsic glory.

When an apple in Eden looked red to us, the apple was gloriously, perfectly, and primitively *red*. There was no need for a long causal chain from the microphysics of the surface through air and brain to a contingently connected visual experience. Rather, the perfect redness of the apple was simply revealed to us. The qualitative redness in our experience derived entirely from the presentation of perfect redness in the world.

Eden was a world of perfect color. But then there was a Fall.

First, we ate from the Tree of Illusion. After this, objects sometimes seemed to have different colors and shapes at different times, even though there was reason to believe that the object itself had not changed. So the connection between visual experience and the world became contingent: we could no longer accept that visual experience always revealed the world exactly as it is.

Second, we ate from the Tree of Science. After this, we found that when we see an object, there is always a causal chain involving the transmission of light from the object to the

Published in (T. Gendler & J. Hawthorne, eds) *Perceptual Experience*. Oxford University Press, 2006. I owe a special debt to George Bealer for suggesting the central metaphor of this chapter, in a memorable conversation in Chicago airport. In discussing whether perfect properties could be instantiated in any possible world, George said “Maybe that’s how it was in Eden.” I also owe a debt to Brad Thompson, whose exploration of a Fregean approach to phenomenal content in his dissertation helped to spark the line of thinking here. Conversations with John Hawthorne and Mark Johnston about primitivist views have also had a significant influence on the ideas here. I first presented this material at a conference on the ontology of color in Fribourg in November 2003, and comments from all of the participants there, especially Alex Byrne, Larry Hardin, Barry Maund, and Martine Nida-Rümelin, were very helpful. Thanks also to audiences at the Universities of Nottingham, Virginia, and Arizona, and at conferences in Santa Barbara and in Florence. For their commentaries at these two conferences, I am grateful to Gideon Rosen, Susanna Siegel, and Aaron Zimmerman.

retina, and the transmission of electrical activity from the retina to the brain. This chain was triggered by microphysical properties whose connection to the qualities of our experience seemed entirely contingent. So there was no longer reason to believe in acquaintance with the glorious primitive properties of Eden, and there was no good reason to believe that objects in the world had these properties at all.

We no longer live in Eden. Perhaps Eden never existed, and perhaps it could not have existed. But Eden still plays a powerful role in our perceptual experience of the world. At some level, perception represents our world as an Edenic world, populated by perfect colors and shapes, with objects and properties that are revealed to us directly. And even though we have fallen from Eden, Eden still acts as a sort of ideal that regulates the content of our perceptual experience. Or so I will argue.

2 Phenomenal Content

My project in this chapter concerns the *phenomenal content* of perceptual experience. This notion can be defined in terms of the notions of *phenomenal character* and *representational content*.

The phenomenal character of a perceptual experience is what it is like to have that experience. Two perceptual experiences share their phenomenal character if what it is like to have one is the same as what it is like to have the other. We can say that in such a case, the experiences instantiate the same phenomenal properties. As I use the term “perceptual experience”, it is true by definition that any perceptual experience has phenomenal character. As I use the term, it is not true by definition that every perceptual experience has an object in the external world: hallucinatory experiences count as perceptual experiences.

A representational content of a perceptual experience is a condition of satisfaction of the experience. I will take it for granted that perceptual experiences can be veridical or falsidical: they can represent the world correctly or incorrectly. Intuitively, perceptual illusions and hallucinations are falsidical, while non-illusory perceptual experiences of objects in the external world are veridical. And intuitively, a given experience will be either veridical or falsidical, depending on what the world is like. If so, we can say that an experience is associated with a condition of satisfaction. If and only if the world satisfies the condition, then experience will be veridical. For example, one might plausibly hold that an ordinary experience of a red square in front of one will be veridical roughly when there is a red square in front of one.

A *phenomenal content* of a perceptual experience is a representational content that is determined by the experience's phenomenal character. More precisely: a representational content C of a perceptual experience E is a phenomenal content if and only if necessarily, any experience with the phenomenal character of E has representational content C.

Put this way, it is a substantive thesis that perceptual experiences have phenomenal content. But there is good reason to believe that they do. The basic reason has been articulated at length by Siewert (1998). It is plausible that perceptual experiences are assessable for accuracy, in virtue of their phenomenal character. Intuitively, by virtue of their phenomenal character, experiences present the world as being a certain way. My experience of a red square in front of me has a certain phenomenal character, and by virtue of this phenomenal character, the experience places a constraint on the world. The world can be such as to satisfy the constraint imposed by the phenomenal character of the experience, or such as to fail to satisfy the constraint. This is to say that the phenomenal character determines a condition of satisfaction for the experience, one that is shared by any experience with the same phenomenal character. This condition of satisfaction will be a phenomenal content.

The plausible thesis that perceptual experiences have phenomenal content leaves many other questions open. For example, the thesis is neutral on whether phenomenal character is prior (in some sense) to representational content, or vice versa. It is compatible with the thesis that phenomenal character is grounded in representational content (as held by Dretske 1995 and Tye 1995, among others), and it is compatible with the thesis that representational content is grounded in phenomenal character (as held by Searle 1990 and Horgan and Tienson 2002, among others).

The thesis also leaves open the nature of phenomenal content. On the face of it, there are many ways to associate representational contents with perceptual experience. For example, one might associate a perceptual experience with an object-involving content (the content that O is F, where O is the object of the experience), an existential property-involving content (for example, the content that there exists something in location L that is F), a content involving modes of presentation of these objects and properties (more on this shortly), and perhaps others.

My own view is that one should be a pluralist about representational content. It may be that experiences can be associated with contents of many different sorts by different relations: we can call such relations *content relations*. For example, there may be one content relation that associates experiences with object-involving contents, and another that associates experiences with existential contents. Each of the different sorts of content, and the corresponding content relations, may have a role to play for different explanatory purposes.

On this view, there may not be such a thing as *the* representational content of a perceptual experience. Instead, a given experience may be associated with multiple representational contents via different content relations.

On the other hand, not all of these contents are equally plausible candidates to be *phenomenal* contents. Some of these contents seem to be such that they can vary independently of the phenomenal character of an experience. If so, they may be representational contents of the experience, but they are not phenomenal contents. More precisely, we can say that a given content relation is a *phenomenal content relation* when any two experiences with the same phenomenal character are related by the content relation to the same content. A phenomenal content relation relates a given experience to a phenomenal content of the experience. For ease of usage, I will speak of *the* phenomenal content of an experience, but we should leave open the possibility that there is more than one phenomenal content relation, so that a given experience can be associated with phenomenal contents of more than one sort. Later in the I will explore this possibility in detail.

In this chapter, I will focus on the question: what is the phenomenal content of a perceptual experience? This is a more constrained question than the corresponding question simply about representational content. But it is an important question to answer. On the face of it, the phenomenal content of an experience is an extremely important aspect of its representational content: it captures a way in which the world is presented in the *phenomenology* of the experience. One can reasonably expect that if we can understand phenomenal content, this will help us to understand the relationship between phenomenal character and representational content, and it may well help us to understand the nature of phenomenology itself.

In what follows, I will first consider and reject one hypothesis about phenomenal content (a Russellian hypothesis) and will argue for another hypothesis (a Fregean hypothesis). In doing so I will cover ground also covered in Chalmers 2004a, although in a slightly different way, so those familiar with that paper might just skim the next two sections. I will then raise problems for the Fregean hypothesis, involving its phenomenological adequacy. I will argue that these problems are best handled by moving to a more refined view of phenomenal content, one that gives Edenic representation a key role.

3 Russellian Content

I will focus on the phenomenal content of visual experiences, and especially of experiences of color. I will take the canonical sort of color experience to be an experience as

of an object having a certain color at a certain location. Our experiences typically present objects to us as having a certain distribution of colors at different locations on their surfaces. A book might be presented to me as being certain shades of blue at some points on its surface and as being certain shades of red at other points. For simplicity, I will focus just on the experience of color at a specific point: for example, the experience of a book's being a specific shade of blue at a specific location on its surface. The conclusions generalize, however, and I will discuss the generalization later in the chapter.

What is the phenomenal content of such a color? That is, what sort of representational content is shared by all experiences with the same phenomenal character as the original experience?

A first attempt at an answer to this question might be the following: the experience represents object O (the book) as having color C (a specific shade of blue) at location L (a particular point in space). So one can associate the experience with the following condition of satisfaction: the experience is satisfied iff object O has color C at location L. This is an *object-involving* content: it involves a specific object O, and its satisfaction depends on the properties of O. We can say that in this case, the experience attributes a certain sort of color property to the object O.

It is plausible that experiences have contents of this sort. When a certain book appears red to us, there is a quite reasonable sense in which the experience will be satisfied iff the book in question is red at the relevant location. That is, satisfaction conditions for experiences can often be understood in terms of the instantiation of certain properties (such as redness) by certain objects (such as the book). In these cases, we can say that the experience *attributes* the property to the object.

It is implausible that this object-involving content is a phenomenal content, however. On the face of it, there might be an experience with the same phenomenal character as the original experience, directed at a quite different object O¹ (perhaps an experience that I could have when looking at a different copy of the same book, for example). And plausibly, there might be an experience with the same phenomenal character as the original experience, directed at no object at all (a hallucinatory experience, for example). These experiences could not have an object-involving content involving the original object O: the former experience will at best have a content involving a different object O¹, and the latter experience may have

no object-involving content at all. If so, the object-involving content is not phenomenal content.¹

A somewhat more plausible candidate for the phenomenal content of the experience is the following: the experience represents that there is an object that has color C at location L. This content is not object-involving: in effect it is existentially quantified, so it does not build in any specific object. Unlike the object-involving content, this content can be possessed by experiences that are directed at different objects or at no object at all. The content is property-involving, however: in effect, it has a color property and a location property as its constituents. For the content to be a remotely viable candidate for phenomenal content, the location property cannot be an *absolute* location property: a phenomenally identical experience might be instantiated in a quite different location, and it is not plausible that this experience could attribute the same absolute location property as the original experience. Rather, the location property must be a *relative* location property: the property of being a certain distance in front of the perceiver at a certain angle, for example.

The contents discussed above are all Russellian contents, in that they are composed from objects and properties. The object-involving content can be seen as a certain structured complex of an object, a location property, and a color property, while the existential content can be seen as a structured complex involving an existential quantifier, a location property, and a color property. We can say that both of these contents involve the attribution of certain specific properties, although in one case the properties are attributed to specific objects, and in the other case to an unspecified object under an existential quantifier. Contents of this sort contrast with Fregean contents, composed from modes of presentation of objects and properties in the world, to be discussed shortly.

¹ One sort of disjunctivism about perceptual experience (disjunctivism about phenomenology) denies that experiences directed at different objects could have the same phenomenology, and denies that a hallucinatory experience could have the same phenomenology as an experience of an external object. On this view, phenomenal content might be object-involving. I will assume that this view is false in what follows. Another variety of disjunctivism (disjunctivism about metaphysics) allows that a hallucinatory experience and an ordinary perceptual experience may share phenomenal character, but holds that they have a fundamentally distinct underlying metaphysical nature: one experience involves an object and one does not. I will be neutral about this view in what follows. A third variety (disjunctivism about content) allows that these experiences have the same phenomenology, but holds that they have different representational contents: for example, experiences of different objects will have different object-involving contents. Proponents of this view will agree that these object-involving contents are not phenomenal contents. They may deny that there is any phenomenal content, holding that the relevant experiences share no content, or they may accept that there is phenomenal content while holding that it is less fundamental than object-involving content. I am inclined to reject both varieties of disjunctivism about content, but I will not assume that they are false. In what follows I will in effect argue against the former view, but some of what I say may be compatible with the latter view.

Let us say that the *Russellian hypothesis* holds that the phenomenal content of a perceptual experience is a sort of Russellian content. To assess this hypothesis, I will henceforth abstract away from issues involving the representation of objects, and will focus on the representation of properties. In particular, I will focus on the representation of color properties—or at least on the representation of properties by color experience. Later, I will discuss the representation of other properties.

On the Russellian hypothesis, the phenomenal content that is distinctively associated with color experience will be a Russellian property-involving content, involving the attribution of a property C. The Russellian hypothesis requires the *Russellian constraint*: all phenomenally identical color experiences attribute the same property to their object.

Strictly speaking, the Russellian hypothesis requires only that *globally* identical experiences attribute the same property to their objects, so that the property attributed may depend on the holistic character of a visual experience, including its spatial phenomenal character and its overall pattern of color phenomenal character. For simplicity, I will usually assume a view on which the property attributed depends on the *local* phenomenal character of the experience, so that any two experiences with the same local phenomenal character in respect of color will attribute the same property. Nothing important will depend on this assumption, however.

We can say two experiences that share their local phenomenal character instantiate the same local phenomenal properties. Local phenomenal properties include properties such as *phenomenal redness*: this is a property instantiated by all experiences that share a certain specific and determinate local phenomenal character, one that is often caused (in us) by seeing things with a certain specific shade of red. (“Phenomenal red₃₁-ness” might be a more apt label for a determinate property than “phenomenal redness”, but I will use expressions of the latter sort for ease of usage.) Note that phenomenal redness (a property of experiences) will plausibly be distinct from ordinary redness, and from the property attributed by phenomenally red experiences.

If the phenomenal content of color experience is Russellian, what sort of properties does it attribute? Intuitively, these properties are color properties: a phenomenally red experience plausibly attributes redness, for example. I will not assume this in what follows, so that room is left open for views on which color experiences represent properties other than color properties. However, the natural hypotheses concerning the nature of the attributed properties correspond fairly closely to the standard range of options concerning the nature of color properties.

One might hold that the properties attributed are *physical* properties: something along the lines of a surface spectral reflectance. One might hold that the properties attributed are *dispositional* properties, involving the disposition to cause a certain sort of experience in appropriate conditions. One might hold that the properties attributed are *mental* properties of some sort: perhaps properties that are actually instantiated by one's experiences or by one's visual fields. Or one might hold that the properties instantiated are *primitive* properties: simple intrinsic qualities, of the sort that might have been instantiated in Eden.

Each of these views is a version of the Russellian hypothesis: we might call them *physicalist*, *dispositionalist*, *projectivist*, and *primitivist* versions of Russellianism concerning phenomenal content. (The projectivist and primitivist views correspond to what Shoemaker (1990) calls "literal" and "figurative" projectivism, except that I take primitivism to be neutral on whether the relevant properties are actually instantiated.) Each of these views is held by some philosophers. For example, the physicalist view is held by Tye (1995); the dispositionalist view is held by Shoemaker (2001); the projectivist view is held by Boghossian and Velleman (1989); and the primitivist view is held by Maund (1995).²

Each of these views has well-known problems. The physicalist view is incompatible with intuitions about spectrum inversion, according to which phenomenally identical color experiences could have represented quite different physical properties in different environments, and it also seems to be incompatible with the internalist intuition that phenomenal character does not constitutively depend on an individual's environment. The dispositionalist view is incompatible with the intuition that color experience attributes nonrelational properties, and also has serious difficulties in individuating the relevant dispositions so that phenomenally identical experiences always attribute the same dispositions. The projectivist and primitivist views suffer from the problem that it seems that the relevant properties are not actually instantiated by external objects, with the consequence that all color experience is illusory.

These problems have been discussed extensively elsewhere by others and by me (see Chalmers 2004a), so I will not dwell on them here. But one can summarize the problems with the following general argument against a Russellian view:

² The physicalist view is also held by Byrne and Hilbert (2003), Dretske (1995), and Lycan (1996). Versions of the dispositionalist view are also held by Egan (forthcoming) and Kriegel (2002). Versions of the primitivist view are also held by Campbell (1993), Holman (2002), Johnston (forthcoming), McGinn (1996), Thau (2002), and Wright (2003). See Stoljar (forthcoming) and Chalmers (2004a) for more discussion of these alternatives.

- (1) Some phenomenally red experiences of ordinary objects are veridical.
- (2) Necessarily, a phenomenally red experience of an object is veridical iff its object instantiates the property attributed by the experience.
- (3) The properties attributed by color experiences are nonrelational properties.
- (4) For any veridical phenomenally red experience of an ordinary object, it is possible that there is a falsidical phenomenally red experience of an object with the same nonrelational properties as the original object.

-
- (5) There is no property that is attributed by all possible phenomenally red experiences.

Here, premise (1) has obvious plausibility, and premise (2) is a natural part of any view on which experiences attribute properties. Premise (3) is grounded in the phenomenology of color experience, and premise (4) corresponds to a fairly weak inversion claim. In fact a version of (4) with a mere existential quantifier instead of a universal quantifier would suffice for the conclusion. But the universally quantified claim seems no less plausible. Take any veridical phenomenally red experience of an ordinary object: say, of an apple. Then it is plausible that there could be a community (one with a somewhat different visual apparatus, or with a somewhat different environment) in which apples of that sort normally cause phenomenally green experiences rather than phenomenally red experiences, and in which phenomenally red experiences are caused by objects of a quite different sort. In such a community, it might happen that on one occasion in unusual conditions, such an apple causes a phenomenally red experience. It is plausible that such an experience would not be veridical.

The conclusion follows straightforwardly from the premises: (1) and (4) entail that there are possible veridical and nonveridical phenomenally red experiences of objects with the same nonrelational properties. Conjoined with (2) and (3), one can conclude that there is no property attributed by both experiences. The conclusion entails that there is no Russellian content that is shared by all phenomenally red experiences (given the natural assumption that if an experience has Russellian content, the property attributed by the experience is attributed by its Russellian content). The argument generalizes straightforwardly from phenomenal redness to any phenomenal property that can be possessed by veridical color experiences, including global phenomenal properties as well as local phenomenal properties. It follows that no veridical experience of an ordinary object has a Russellian phenomenal content. Given that

there are veridical experiences of ordinary objects (premise (1)), it follows that the Russellian hypothesis about color experience is false.

One can strengthen the argument by noting that there is no requirement that the original and inverted communities perceive distinct apples. If we replace premises (1) and (4) by the premise that there can be phenomenally identical veridical and nonveridical experiences of the same object, then even without premise (3), it follows that there is no property that is attributed by all phenomenally red experiences. For there is no property that the apple simultaneously possesses and lacks: not even a relational or dispositional property.

(A possibility left open is that what is attributed is what we might call a “relational property radical”: perhaps something like *normally causes phenomenally red experiences in ---*, where the open place is to be filled by the subject of the experience (see Egan (forthcoming) for a view in this vicinity). We could say that different subjects attribute the same relational property radical, which determine different relational properties for different subjects. Of course relational property radicals are not really properties, and this proposal is also subject to the usual phenomenological objections.)

Proponents of Russellian views will respond by denying one of the premises of the argument: depending on the view, they might deny premise (1), (3), or (4) of the original argument. This leads to a well-worn dialectic that I do not want to get into here. For now I will just note that each of the premises enjoys strong intuitive support, and I will take this as strong *prima facie* reason to believe that the Russellian hypothesis is false.

If we accept premises (1)-(4), we are left with something like the following view. Color experiences attribute nonrelational properties, and these properties are sometimes instantiated by ordinary objects. The most plausible candidates for such properties are intrinsic physical properties, such as surface reflectance properties³: so our phenomenally red experiences might attribute a specific physical property, which we can call physical redness. On this view, color experiences have Russellian content that involve the attribution of these properties. However, it is possible that experiences with the same phenomenal character (perhaps in a different community) can have different Russellian contents of this sort, due to differences in the environment of the perceivers. So Russellian content is not phenomenal content.

³ One might understand a surface reflectance property as a sort of dispositional property, involving dispositions to reflect certain sorts of light. If so, then the physical properties attributed by color experiences should probably be understood to be the categorical bases of surface reflectance properties, rather than reflectance properties themselves. I will usually simply talk of reflectance properties in what follows, however.

4 Fregean Content

If one thought that all content were Russellian content, one might conclude from the above that experiences (or at least color experiences) do not have phenomenal content. But this would require denying the strong intuition that experiences are assessable for accuracy in virtue of their phenomenal character. The alternative is to hold that phenomenal content is something other than Russellian content. There is a natural alternative: Fregean content.

Where Russellian content involves objects and properties, Fregean content involves modes of presentation of objects and properties. The idea is familiar from the philosophy of language. Utterances of “Hesperus is a planet” and “Phosphorus is a planet” have the same Russellian content (a content attributing planethood to Venus), but a Fregean view holds that they have different Fregean contents. The terms “Hesperus” and “Phosphorus” are associated with different modes of presentation of the planet Venus (different Fregean senses, for example), mirroring the difference in their associated cognitive roles. This difference in modes of presentation makes for a different Fregean content for the two utterances.

A natural view of Fregean modes of presentations holds that they are *conditions on extension*. The extension of an expression is something like an object or a property. The mode of presentation associated with an expression is a condition that an object or property must satisfy in order to qualify as the expression’s extension. For a term such as “Hesperus”, for example, one might plausibly hold that the associated mode of presentation is a condition that picks out the bright object at a certain position in the evening sky. For “Phosphorus”, the mode of presentation might be a condition that picks out the bright object at a certain position in the morning sky. In the actual world, both conditions are satisfied by the planet Venus, so Venus is the extension of both expressions.

A pluralist view can allow that utterances have both Russellian and Fregean content, under two distinct content relations. An utterance of “Hesperus is Phosphorus” might be associated with both a Russellian content (holding that Venus is identical to itself) and a Fregean content (holding roughly that the bright object in the morning sky is identical to the bright object in the evening sky). Both of these contents (and the corresponding content relations) may be useful for different explanatory purposes: the former may be more relevant to evaluating the sentence in counterfactual circumstances, for example, while the latter may be more relevant to analyzing the sentence’s epistemic role.

This model also applies to terms that refer to perceptible properties such as colors. For example, a color term such as “red” might refer to the same physical property as a physical term “P”. Here, the two expressions will be associated with quite different cognitive roles: an

utterance of “red is P” will be a cognitively significant claim. So there is reason to believe that the color term “red” has a Fregean content quite distinct from that of the physical expression “P”, even though the two expressions have the same Russellian content.

At this point, it is natural to suggest that perceptual experiences, as well as linguistic expressions, have both Russellian and Fregean content. The Russellian content of a color experience will be the property attributed by the experience. If premises (1)-(3) of the earlier argument are correct, this will plausibly be an intrinsic physical property, such as a surface reflectance property: a phenomenally red experience may attribute physical redness, for example. The Fregean content will be a mode of presentation of that property. On the face of it, color experiences attribute colors under a distinctive mode of presentation (one quite distinct from a physical mode of presentation of the color, for example, and much closer to the mode of presentation associated with a color term). It is natural to suggest that this distinctive mode of presentation corresponds to a distinctive sort of Fregean content. This allows the attractive suggestion that where phenomenally identical experiences can have different Russellian contents, they will always have the same Fregean content. If so, Fregean content is phenomenal content.

What is the Fregean content associated with a color experience? It is a condition that a property must satisfy in order to be the property attributed by the experience. There is a natural candidate for such a condition. Let us assume, as before, that the property attributed by a phenomenally red experience is a physical property such as physical redness. Then one can naturally hold that the associated condition on this property is the following: it must be the property that normally causes phenomenally red experiences (in normal conditions for the perceiver). Plausibly, it is precisely because physical redness satisfies this condition that it is the property attributed by phenomenally red experiences. If so, this suggests that the Fregean content of such an experience is precisely this condition.

This Fregean content is a natural candidate for the phenomenal content associated with phenomenal redness. Certainly this content accommodates inversion scenarios quite straightforwardly. In an environment where phenomenally red experiences are normally caused by physically green things, such an experience will attribute physical greenness, so that its Russellian content differs from the Russellian content of a phenomenally red experience in our environment. But its Fregean content will be exactly the same: both are the condition that picks the property that normally causes phenomenally red experiences in the perceiver. (Here the perceiver will be picked out under an indexical mode of presentation that can be shared between two different perceivers; see the appendix for a way to model this using centered worlds.) Due to differences in the environment, this common Fregean content

yields distinct Russellian contents: the condition picks out physical redness in an ordinary environment, and physical greenness in the alternative environment. All this suggests that Fregean content is a plausible candidate to be phenomenal content.

The Fregean content of a given color experience can itself be seen as a condition of satisfaction. To a first approximation, the Fregean content of a phenomenally red experience will be satisfied when there is an object at the appropriate location relative to the perceiver that instantiates the property that normally causes phenomenally red experiences in the perceiver. To a second approximation, one might want to give a corresponding Fregean treatment to the attributed location property (as I will discuss later), and one might want to give a Fregean treatment to the object of the experience (for example, holding that the Russellian content of the experience is the specific object, and that the corresponding Fregean content is the condition that picks out the object that is causing the current experience). For now, I will abstract away from these matters and concentrate on that aspect of the content that is associated with color. But however we flesh out the details, this Fregean content will be a condition of satisfaction for the experience. If the perceiver's environment meets the condition, then the experience will be veridical; if it does not, the experience will be nonveridical.

The Fregean view of phenomenal content can be seen as combining aspects of a physicalist and dispositionalist view of the content of color experiences. As on the physicalist view, physical properties (such as surface properties) constitute the Russellian content of color experience, although on the Fregean view this content is not phenomenal content. And as on the dispositionalist view, dispositions to cause certain sorts of experiences are central to the phenomenal content of color experience, although on the Fregean view this content is not Russellian content. Instead, the disposition serves as a sort of Fregean mode of presentation for the physical property that is attributed by the experience.

Compared with the Russellian physicalist view of phenomenal content, the Fregean view has the advantage that it can accommodate inversion scenarios straightforwardly. It can also straightforwardly reconcile the environment dependence of Russellian content with the environment-independence of phenomenology. There is good reason to think that Fregean content is a sort of *narrow* content, which depends only on the internal state of the individual and not on the environment (Chalmers 2002a). This allows us to combine the view that phenomenology is internally determined with the view that phenomenology is intrinsically representational. The sort of content that is intrinsic to phenomenology is (environment-independent) Fregean content. When situated in a particular environment, this yields a Russellian content.

Compared with the Russellian dispositionalist view, the Fregean view has the advantage that it can accommodate the intuition that experiences attribute nonrelational properties. Furthermore, putting the disposition in the mode of presentation removes worries about the individuation of the relevant dispositions: the idea of an indexical property is obscure, but there are already good reasons to believe in indexical modes of presentation. Finally and importantly, the Fregean view can accommodate the strong intuition that things could have been as they perceptually seem to be, even had there been no observers. Counterfactual judgments of this sort generally reflect Russellian content rather than Fregean contents (more generally, they reflect the second dimension in the two-dimensional framework rather than the first). So a view on which Russellian contents are dispositional does not respect these judgments, but a view on which Russellian contents are physical properties (picked out under a dispositional mode of presentation) delivers the intuitively correct results.

5 Phenomenological Adequacy

The hypothesis that phenomenal content is Fregean content has many virtues. In particular, it seems to capture our intuitions about the environments in which an experience with a given phenomenal character will be veridical, yielding a condition of satisfaction that is determined by phenomenal character. Still, there remains a cluster of worries about the view.

This cluster of worries concerns what we might call the *phenomenological adequacy* of the view. Simply put, the worry is that Fregean content does not seem to adequately reflect the phenomenal character of an experience. In particular, one can argue that when we introspect and reflect on the way that the world is presented in the phenomenology of perceptual experience, the phenomenology seems to have properties that are in tension with the Fregean view of phenomenal content. These properties include the following:

Relationality: Intuitively, it seems to us that when we have an experience as of a colored object, there is a certain property (intuitively, a color property) that the object seems to have. And intuitively, it is natural to hold that the phenomenology of the experience alone suffices for it to seem that there is an object with that very property. That is, reflection on phenomenology suggests that there is an internal connection between phenomenology and certain properties that objects seem to have. One could summarize this by saying that the phenomenology of color experience seems to be *relational*: by virtue of its phenomenology, a specific color experience seems to relate us to a specific color property. If this point is correct, it suggests that color experiences have Russellian phenomenal content.

In a critical discussion of the Fregean view, Shoemaker (forthcoming) brings this point out by an appeal to the Moorean “transparency” intuition. According to this intuition, we attend to the phenomenal character of an experience by attending to the properties that objects in the world appear to have. An extension of this intuition suggests that we discern similarities and differences in phenomenal character by discerning similarities and differences in the properties that objects in the world seem to have. This suggests a strong connection between phenomenal character and Russellian content. Shoemaker says:

The phenomenal character of veridical experiences of a given color can be different in different circumstances (e.g., different lighting conditions), and for creatures with different sorts of perceptual systems. So the same color will have to have a number of distinct modes of presentation associated with it. To say that this variation is only a variation in the *how* of perceptual representation and in no way a variation in *what* is represented, seems to me at odds with the phenomenology. When the light brown object in shadow and the dark brown object not in shadow look the same to me, the sameness is experienced as being *out there*—and in such a case the perception can be perfectly veridical. Similarity in the presenting manifests itself in represented similarity in what is represented, and in the absence of perceptual illusion requires that there is similarity in what is represented. More generally, the best gloss on the Moorean transparency intuition is that the qualitative character that figures in the perception of the color of an object is experienced as in or on the perceived object. (Shoemaker, forthcoming)

One can also bring out the point by appealing to an inversion scenario. Jack and Jill are phenomenal duplicates, but live in different environments. Jack’s phenomenally green experiences are normally caused by objects with property X, while Jill’s experiences are normally caused by objects with property Y. Shoemaker’s point suggests that even if Jack’s and Jill’s experiences are associated with distinct properties (X and Y), there is a strong intuitive sense in which the objects look to be the same to Jack and to Jill. That is, the phenomenal similarity suggests that there is a common property (intuitively, a sort of greenness) such that the relevant objects look to have that property both to Jack and to Jill.

This intuitive point stands in tension with the Fregean view. The Fregean view entails that Jack’s and Jill’s experiences share a mode of presentation, but it does not entail that the experiences represent a common property. In fact, it suggests that Jack’s and Jill’s experiences represent distinct properties, X and Y. So it is difficult for the Fregean view to accommodate any internal connection between an experience’s phenomenal character and the properties that it represents.

A related point is that phenomenologically, a color experience appears to represent an object as having a certain *specific* and *determinate* property. Intuitively, this specificity and

determinacy is tied very closely to the specific and determinate phenomenal character of the experience. According to the Fregean view, while an experience may represent a specific and determinate property, its phenomenal character leaves the nature of this property wide open: the determinate property represented may depend on matters quite extrinsic to the phenomenology. This seems to conflict with a strong phenomenological intuition.

Simplicity: A second objection is that Fregean contents seem to be overly complex: one might say that they “overintellectualize” the content of an experience. According to this objection, the phenomenological structure of a visual experience is relatively simple: it represents certain objects as having color and shape properties, and so on. But one cannot find anything like “the normal cause of such-and-such experience” in the visual phenomenology. On the face of it, the “normal cause” relation is not phenomenologically present at all: it is something imposed after the fact by theorists, rather than directly reflecting the experience’s phenomenology.

A related objection turns on the fact that Fregean contents require reference to *experiences*: properties are picked out as the normal cause of a certain type of experience, and objects might be picked out as the cause of a certain token experience. But one can object that the perceptual phenomenology does not (or at least, need not) involve representation of experiences: it need only involve representation of the world. This is another often-invoked aspect of the “transparency” of experiences: the phenomenology of perception usually seems to present the world directly, not in virtue of representation of any experiential intermediaries. Again, to invoke the representation of experience seems to overintellectualize the experience, by introducing complexity that is not apparent in the experience’s phenomenology.

Internal unity: A final objection is that it seems that there can be internal unity among the contents of experiences that have quite different phenomenal character. For example, one can argue that there is an internal unity between the representation of space in visual and tactile experience, by virtue of which these are constrained to represent a common set of spatial properties. Phenomenologically, it seems that when an object looks flat and when it feels flat, it looks and feels to have the same property (flatness). This commonality seems to hold by virtue of an internal relationship between the phenomenology of visual and tactile experiences. It is arguable that something similar applies to experiences as of the same color in quite different lighting conditions. For example, experiences of a white object in shadow and out of shadow may have quite different phenomenal characters, but it is arguable that the experiences are internally related in a way so that both represent the object as being white.

This internal unity is not straightforwardly accommodated by a Fregean view (assuming that the Fregean view might also apply to experiences of space). One might think that because

visual and tactile experiences of space are phenomenally quite different, they will be associated with quite different Fregean modes of presentation. One will represent the normal cause of certain visual experiences, and another will represent the normal cause of certain tactile experiences. It might turn out that as a matter of contingent fact these normal causes coincide, so that the properties represented coincide, but nothing in the experiences themselves guarantees this. This stands in tension with the intuition above that there is an internal phenomenological connection between the tactile and visual representation of space, according to which these have common contents by virtue of their phenomenology. The same goes for the case of phenomenally different experiences as of the same color: the Fregean view suggests that these will have distinct modes of presentation that at best contingently pick out a common property, which stands in tension with the intuition that these experiences have common representational content by virtue of their phenomenology.

I do not think that any of these three objections—from relationality, simplicity, and internal unity—are knockdown objections to the Fregean view. For a start, all of them rest on phenomenological intuitions that could be disputed. I will not dispute them, however: I am inclined to give each of the intuitions some *prima facie* weight. But even if one takes the intuitions at face value, it is not clear that any of them entail that the Fregean view is false. Rather, I think all of them can be seen as pointing to a certain *incompleteness* in the Fregean view: the Fregean account so far is not a full story about the phenomenal content of experience. For a full story, the Fregean view needs to be supplemented.

The relationality objection, for example, suggests that there is a Russellian aspect to the phenomenal content of perceptual experience: that phenomenally identical experiences involve representation of some common property. The intuitions here are somewhat equivocal: in the Jack and Jill case, for example, at the same time as we have the intuition that some common property is phenomenologically represented (as a Russellian view of phenomenal content would suggest), we also have the intuition that different properties might be represented by virtue of distinct environmental connections (as a Fregean view of phenomenal content would suggest). If we are pluralists about content, these two intuitions need not contradict each other. Rather, they might be reconciled if we adopt a view on which there is both a Russellian and a Fregean aspect to the phenomenal content of experiences. The intuition here does not entail that Fregean content is not phenomenal content: rather, it suggests that Fregean content is not all there is to phenomenal content.

The force of the simplicity objection is somewhat unclear. Construed as an argument against Fregean phenomenal content, it turns on the tacit premise that the phenomenal content of an experience must have a structure that directly mirrors the phenomenological structure of

the experience (or perhaps that it directly mirrors the way it seems to us on introspection that the world is perceptually presented). We might call this somewhat elusive idea the “mirroring constraint”. A proponent of the Fregean view might reply simply that the mirroring constraint is an unreasonable constraint on an account of the phenomenal content of experience. As we have defined it, phenomenal content is content that supervenes on the phenomenal character of an experience, but there is nothing in this definition that requires a tighter connection than mere supervenience. And the simplicity objection does not give any reason to deny supervenience. So the Fregean may hold that unless one has an argument that supervenience of content on phenomenal character requires mirroring (or unless we redefine the notion of phenomenal content to build in the mirroring constraint), there is no objection to the claim that Fregean content is phenomenal content.

Still, the simplicity objection once again suggests a certain incompleteness in the Fregean view. One might reasonably hold that the supervenience of content on phenomenal character requires some sort of explanation. If there were a direct correspondence between the elements of the content and the elements of phenomenal character, this explanation would be much easier to give. As it is, the extra complexity of Fregean content (such as the invocation of causation and of experience) raises the question of how this complex content is connected to the simple experience. In particular, if one adopts a view on which phenomenal content is somehow grounded in the phenomenology of an experience, then one will need to tell a story about how a complex Fregean content can be grounded in a simple experience. And if one thinks that the phenomenology of an experience is grounded in its phenomenal content, then the same applies in reverse. So there is at least a significant explanatory question here.

Finally, the Fregean view could handle the internal unity objection by saying that visual and tactile experiences of space share a common phenomenal type (in effect, a crossmodal type), and it is this phenomenal type that is relevant to the Fregean mode of presentation of these experiences (“the property that normally causes experiences of type T”). If so, then the different experiences will be constrained to represent a common class of properties. One could likewise suggest that phenomenally distinct experiences of the same color (shadowed and unshadowed, for example) share a phenomenal type, with the same conclusion. This raises the question, however, of just how we assign the relevant phenomenal types. Any given experience belongs to many different phenomenal types, and the selection of the crossmodal phenomenal type (in the spatial case) or the phenomenal type shared by shadowed and unshadowed experiences (in the color case) may seem suspiciously ad hoc. At least, we need to fill in the Fregean view with an account of how the mode of presentation associated with a

given experience is determined, by specifying a principled basis for the choice of a phenomenal type.

One can summarize the worries above by saying that as it stands, the Fregean view does not seem to fully reflect the *presentational phenomenology* of perceptual experience: the way that it seems to directly and immediately present certain objects and properties in the world. It is natural to hold that this presentational phenomenology is closely connected to the phenomenal content of experience. So, to make progress, we need to attend more closely to this presentational phenomenology, and to how it might be connected to phenomenal content.

6 Back to Primitivism

It is useful at this point to ask: what view of the content of perceptual experience is the most phenomenologically adequate? That is, if we were simply to aim to take the phenomenology of perceptual experience at face value, what account of content would we come up with? In particular, what view of the content of color experience best mirrors its presentational phenomenology?

Here, I think the answer is clear. The view of content that most directly mirrors the phenomenology of color experience is primitivism. Phenomenologically, it seems to us as if visual experience presents simple intrinsic qualities of objects in the world, spread out over the surface of the object. When I have a phenomenally red experience of an object, the object seems to be simply, primitively, *red*. The apparent redness does not seem to be a microphysical property, or a mental property, or a disposition, or an unspecified property that plays an appropriate causal role. Rather, it seems to be a simple qualitative property, with a distinctive sensuous nature. We might call this property perfect redness: the sort of property that might have been instantiated in Eden.

One might say: phenomenologically, it seems that visual experience presents the world to us as an Edenic world. Taking the phenomenology completely at face value, visual experience presents a world where perfect redness and perfect blueness are instantiated on the surface of objects, as they were in Eden. These are simple intrinsic qualities whose nature we seem to grasp fully in perceptual experience. For the world to be *exactly* the way that my phenomenology seems to present it as being, the world would have to be an Edenic world in which these properties are instantiated.

This suggests a view on which color experiences attribute primitive properties such as perfect redness and perfect blueness to objects. On this view, color experiences have a Russellian content involving the attribution of these primitive properties. Furthermore, this

content is naturally taken to be phenomenal content. Intuitively, the nature of the primitive properties that are presented to one is fully determined by the phenomenology of the experience: if an experience attributes a primitive property, any phenomenally identical experience will attribute the same primitive property. So this view is a sort of Russellian primitivism about phenomenal content.

For all its virtues with respect to phenomenological adequacy, the Russellian primitivist view has a familiar problem. There is good reason to believe that the relevant primitive properties are not instantiated in our world. That is, there is good reason to believe that none of the objects we perceive are perfectly red or perfectly green. If this is correct, then the primitivist view entails that all color experiences are illusory.⁴

A first reason for doubting that these properties exist came when we ate from the Tree of Illusion. This made it clear that there is no necessary connection between primitive properties and perceptual experiences, and strongly suggested that if there is a connection, it is merely causal and contingent. And once we have accepted that one sometimes has phenomenally red experiences in the absence of perfect redness, it is natural to start to wonder whether the same goes for all of our phenomenally red experiences. This is a relatively weak reason, as the existence of illusions is compatible with the existence of veridical perception, but it is enough to generate initial doubts.

A second and stronger reason came when we ate from the Tree of Science. Science suggests that when we see a red object, our perception of the object is mediated by the reflection or radiation of light from the surface of the object to our eyes and then to our brains. The properties of the object that are responsible for the reflection or radiation of the light appear to be complex physical properties, such as surface spectral reflectances, ultimately grounded in microphysical configurations. Science does not reveal any primitive properties in the object, and furthermore, the hypothesis that objects have the relevant primitive properties seems quite unnecessary in order to explain color perception.

Still, someone might suggest that objects have the primitive properties all the same, perhaps supervening in some fashion on the microphysical properties of the object. In response, one might suggest that this picture will metaphysically complicate the world. It seems at least conceivable that objects with the relevant microphysical properties could fail to instantiate the relevant primitive properties. So it looks as if the relevant primitive properties are a significant addition to the world, over and above the microphysical

⁴ In practice, primitivists are divided on this issue. For example, Holman, Maund and Wright hold that the primitive properties are uninstantiated and that color experiences are illusory, while Campbell, Johnston, and McGinn hold that primitive properties are instantiated and that color experiences can be veridical.

supervenience base. A primitivist might respond in turn by denying that any metaphysical addition is involved (perhaps denying an inference from conceivability to metaphysical possibility), or by accepting that physicalism about ordinary objects is false.⁵ But even if so, there is a remaining problem.

The third and strongest reason for doubting that primitive properties are instantiated stems from an elaboration of the inversion argument given earlier.⁶ Take an ordinary object, such as a red apple. It is familiar from everyday experience that such an object can cause phenomenally red experiences of the apple and (in some circumstances) can cause phenomenally green experiences of the apple, without any change in its intrinsic properties. It then seems that there is no obstacle to the existence of a community in which objects with the intrinsic properties of this apple *normally* cause phenomenally green experiences. We can even imagine that the very same apple normally causes phenomenally red experiences in one community and normally causes phenomenally green experiences in the other.

We can now ask: when a subject in the first community has a phenomenally red experience of the apple, and a subject in the second community has a phenomenally green experience of the apple, which of these experiences is veridical?

Intuitively, there is a case for saying that both experiences are veridical. But this is an unhappy answer for the primitivist. On the primitivist view, any phenomenally red experience attributes perfect redness, and any phenomenally green experience attributes perfect greenness. If both experiences are veridical, it follows that the apple instantiates both perfect redness and perfect greenness. The argument generalizes: for any phenomenal color, it seems that there is a community in which the apple normally causes experiences with that phenomenal color. Taking the current line, it will follow that the apple instantiates every perfect color! The choice of an apple was perfectly general here, so it seems to follow that every object instantiates every perfect color. It follows that no color experience of an object can be illusory with respect to color. Whatever the phenomenal color of the experience, the object will have the corresponding primitive property, so the experience will be veridical. This conclusion is perhaps even more counterintuitive than the conclusion that all color experiences are illusory.

A primitivist might suggest that one of the experiences is veridical and one of them is not. But this imposes an asymmetry on what otherwise seems to be a quite symmetrical

⁵ Among primitivists who think that the primitive properties are instantiated, Campbell and McGinn suggest that they metaphysically supervene on microphysical or dispositional properties, so that they are not a metaphysical addition in the strong sense, while Johnston seems willing to accept that they are a strong metaphysical addition.

⁶ A version of this sort of argument is deployed by Edwards (1998) against Campbell's version of primitivism.

situation. When a subject in one community has a phenomenally red experience of the apple and a subject in the other community has a phenomenally green experience of the apple, both subjects' perceptual mechanisms are functioning in the way that is normal for those communities. And the perceptual mechanisms themselves, involving light and brain, seem to be symmetrically well-functioning in both communities. Perhaps a primitivist can hold the line and assert that one of the experiences is veridical and one is falsidical, simply because the apple is perfectly red and it is not perfectly green. But this line leads to the conclusion that color experiences in one of the communities are *normally* falsidical (after all, objects like the apple normally cause phenomenally green experiences in that community) where corresponding experiences in the other community are normally veridical.

Apart from the unappealing asymmetry, this view yields a serious skeptical worry: it seems that we have little reason to believe that we are in a community that normally perceives veridically as opposed to falsidically. After all, nature and evolution will be indifferent between the two communities above. Evolutionary processes will be indifferent among perceivers in which apples produce phenomenally red experiences, perceivers in which apples produce phenomenally green experiences, and perceivers in which apples produce phenomenally blue experiences. Any such perceiver could easily come to exist through minor differences in environmental conditions or brain wiring. If we accept the reasoning above, only a very small subset of the class of such possible perceivers will normally have veridical experiences, and there is no particular reason to think that we are among them.

Once these options are ruled out, the reasonable conclusion is that neither experience is veridical: the apple is neither perfectly red nor perfectly green. Generalizing from this case, this reasoning suggests that primitive properties are not instantiated at all. I think that this is clearly the most reasonable view for a primitivist to take: on this view, experiences attribute primitive properties, but their objects never possess these properties.

Still, this view has the consequence that all color experiences are illusory. This is a counterintuitive conclusion, and runs counter to our usual judgments about the veridicality of experience. On the face of it, there is a significant difference between a phenomenally red experience of a red wall and a phenomenally red experience of a white wall that looks red because (unknown to the subject) it is bathed in red light. As we ordinarily classify experiences, the former is veridical and the latter is not. In classifying both experiences as falsidical, primitivism cannot respect this distinction.

7 Perfect and Imperfect Veridicality

Here is where things stand. The Fregean view of phenomenal content seems to most accurately capture our judgments about veridicality, but it is not especially phenomenologically adequate. The primitivist view of phenomenal content is the most phenomenologically adequate view, but it yields implausible consequences about veridicality. For a way forward, what we need is an account that captures both the phenomenological virtues of the primitivist view and the truth-conditional virtues of the Fregean view. In what follows I will argue that such an account is available.

One can begin to motivate such a view with the following pair of intuitions.

1. For a color experience to be *perfectly veridical*—for it to be as veridical as it could be—its object would have to have perfect colors. The perfect veridicality of color experience would require that our world is an Edenic world, in which objects instantiate primitive color properties.

2. Even if the object of an experience lacks perfect colors, a color experience can be *imperfectly veridical*: veridical according to our ordinary standard of veridicality. Even after the fall from Eden, our imperfect world has objects with properties that suffice to make our experiences veridical, by our ordinary standards.

This pair of intuitions is strongly supported, I think. The first is supported by the phenomenological observations in the previous section. If we were to take our experience *completely* at face value, we would accept that we were in a world where primitive properties such as perfect redness and perfect blueness are spread homogeneously over the surface of objects. The second is supported by our ordinary judgments about veridicality. When an ordinary white wall looks white to us, then even if it merely instantiates physical properties and not perfect whiteness, it is good enough to qualify as veridical by our ordinary standards.

These two intuitions need not be taken to contradict each other. Instead, they suggest that we possess two notions of satisfaction for an experience: perfect and imperfect veridicality. An experience can be imperfectly veridical, or veridical in the ordinary sense, without being perfectly veridical.

The terminology should not be taken to suggest that when an experience is imperfectly veridical, it is not really veridical. In fact, it is plausible that imperfect veridicality is the property that our ordinary term “veridicality” denotes. We speak truly when we say that a phenomenally red experience of an ordinary red object is veridical. It is just that the experience is not perfectly veridical. To capture this, one could also call imperfect veridicality “ordinary veridicality”, or “veridicality simpliciter”. Or one could use “veridical” for

imperfect veridicality and “ultraveridical” for perfect veridicality. But I will usually stick to the terminology above.

Corresponding to these distinct notions of satisfaction, one will have distinct associated conditions of satisfaction. Imperfect veridicality will be associated with something like the Fregean condition of satisfaction discussed earlier: a phenomenally red experience will be perfectly veridical iff its object has the property that normally causes phenomenally red experiences. Perfect veridicality will be associated with the primitivist condition of satisfaction: a phenomenally red experience will be perfectly veridical iff its object instantiates perfect redness.

Imperfect and perfect veridicality can therefore be seen as associated with distinct *contents* of an experience. We might call the content associated with perfect veridicality the *Edenic content* of an experience, and the content associated with imperfect veridicality the *ordinary content* of the experience.

As we have already seen, our ordinary assessments of veridicality can be seen as associated with two contents in turn. For example, a phenomenally red experience has a Fregean content (satisfied iff its object has the property that normally causes phenomenally red experience) and a Russellian content (satisfied iff its object has physical redness). We might call these contents the *ordinary Fregean content* and the *ordinary Russellian content* of the experience.

One could also, in principle, associate assessments of perfect veridicality with both a Fregean and a Russellian content. But here the Fregean content is much the same as the Russellian content. The Russellian content involves the attribution of perfect redness: it is satisfied in a world iff the relevant object is perfectly red there. Unlike the ordinary Russellian content above, this content does not depend on how the subject’s environment turns out. Regardless of how the environment turns out, the experience in question will attribute perfect redness. So there is no nontrivial dependence of the property attributed on the way the subject’s environment turns out. It follows that the Edenic Fregean content of the experience (which captures the way that the perfect veridicality of the experience depends on the way the environment turns out) is satisfied iff the object of the experience has perfect redness. There may be some differences between the Edenic Fregean and Russellian contents here in the treatment of objects (as opposed to properties), and in the formal modeling (with worlds and centered worlds), but where the color-property aspect of the content is concerned, the contents

behave in very similar ways. So for most purposes one can simply speak of the Edenic content of the experience, one that is satisfied iff a relevant object has perfect redness.⁷

So we have found three distinctive sorts of content associated with an experience: an Edenic content, an ordinary Fregean content, and an ordinary Russellian content. We have seen already that the ordinary Russellian content is not plausibly a phenomenal content: phenomenally identical experiences can have the same (ordinary) Russellian contents. However, for all we have said, both Edenic contents and ordinary Fregean contents are phenomenal contents. It is plausible that any phenomenally red experience will have the Fregean condition of satisfaction above (where satisfaction is understood as imperfect veridicality) and will have the primitivist condition of satisfaction above (where satisfaction is understood as perfect veridicality). So we have more than one phenomenal content for an experience, depending on the associated notion of satisfaction.

8 A Two-Stage View of Phenomenal Content

Perfect and imperfect veridicality are not independent of each other. It is plausible to suggest that there is an intimate relation between the two, and that there is an intimate relation between the associated sorts of phenomenal content.

A natural picture of this relation suggests itself. A phenomenally red experience is perfectly veridical iff its object instantiates perfect redness. A phenomenally red experience is imperfectly veridical iff its object instantiates a property that *matches* perfect redness. Here, to match perfect redness is (roughly) to play the role that perfect redness plays in Eden. The key role played by perfect redness in Eden is that it normally brings about phenomenally red experiences. So a property matches perfect redness if it normally causes phenomenally red experiences. This yields a condition of satisfaction that mirrors the ordinary Fregean content above.

The notion of matching is what links imperfect veridicality to perfect veridicality. I will say more about this notion later, but one can motivate the idea as follows. For our experiences to be perfectly veridical, we would have to live in Eden. But we have undergone the fall from

⁷ In terms of the two-dimensional framework, one can say that phenomenal color properties (at the standard of perfect veridicality) are associated with the same primary and secondary intension. In this way they are reminiscent of expressions such as “consciousness”, “philosopher”, and “two”, which also arguably have the same primary and secondary intensions. These terms can be seen as “semantically neutral” (Chalmers 2004b), as witnessed by the fact that their content does not seem to have the same sort of dependence on empirical discoveries about the environment as terms such as “water” and “Hesperus”. One might say that perceptual representations of perfect redness are semantically neutral in an analogous way.

Eden: no primitive color properties are instantiated by objects in our world. So the best that objects in our world can do is to have properties that can play the role that primitive properties play in Eden. Of course no property instantiated in our world can play that role perfectly, but some can play it well enough, by virtue of normally bringing about phenomenally red experiences. Such a property might be called *imperfect redness*. In our world, imperfect redness is plausibly some sort of physical property, such as a surface spectral reflectance.

More generally, the following is a plausible thesis. If an experience is such that its perfect veridicality conditions require the instantiation of primitive property X, then the experience's imperfect veridicality conditions will require the instantiation of a property that matches X. As before, a property matches X (roughly) if it plays the role that X plays in Eden. The key role is causing experiences of the appropriate phenomenal type. In our world, these properties will typically be physical properties: the imperfect counterparts of X.

This relation suggests the following *two-stage* picture of the phenomenal content of experience. On this picture, the most fundamental sort of content of an experience is its Edenic content, which requires the instantiation of appropriate primitive properties. This content then determines the ordinary Fregean content of the experience: the experience is imperfectly veridical if its object has properties that match the properties attributed by the experience's Edenic content.⁸

On the two-stage view, the ordinary Fregean content of a phenomenally red experience will be satisfied (in an environment) iff a relevant object instantiates a property that matches perfect redness (in that environment). This ordinary Fregean content will itself be associated with an ordinary Russellian content: one that is satisfied iff the (actual) object of the experience has P, where P is the property that matches perfect redness in the environment of the original experience. On this view, all phenomenally red experiences will have the same Fregean content, but they may have different Russellian contents, depending on their environment.

⁸ Aaron Zimmerman suggested that instead of associating an experience with two contents, we could associate an experience with a single graded content that has degrees of satisfaction: the content might be perfectly satisfied, imperfectly satisfied, and so on, depending on how the world turns out. A pluralist can allow that we can associate experiences with graded contents like these. However, this single graded content will lose some of the structure present in the dual contents: in particular, we cannot easily analyze it in terms of attribution of a property to objects in the environment, and the matching relation between Edenic and ordinary content will not easily be reflected in this account. So this picture will lose some of the explanatory structure that is present on the two-stage view.

Of course this Fregean content gives exactly the same results as the Fregean content discussed earlier: an object will instantiate a property that matches perfect redness iff it instantiates a property that normally causes phenomenally red experiences. But the two-stage view gives a more refined account of how this Fregean content is grounded, one that more clearly shows its roots in the phenomenology of the experience. The view also has the promise of being more phenomenologically adequate than the original Fregean view seemed to be, by giving a major role to the Edenic content that directly reflects the experience's phenomenology. The resulting view is a sort of semi-primitivist Fregeanism: a version of the Fregean view on which the Fregean content is grounded in a primitivist Edenic content.

On this view, Eden acts as a sort of regulative ideal in determining the content of our color experiences. Our world is not Eden, but our perceptual experience requires our world to match Eden as well as possible. Eden is central to the content of our experience: it is directly reflected in the perfect veridicality conditions of the experience, and it plays a key role in determining the ordinary veridicality conditions of our experiences.

One might put the two-stage view as follows: our experience *presents* an Edenic world, and thereby *represents* an ordinary world. We might say that the perfect veridicality conditions of the experience are its *presentational content*, and the imperfect veridicality conditions of the experience are its *representational content*. As pluralists we can allow that experiences have both sorts of content, with an intimate relation between them. Presentational content most directly reflects the phenomenology of an experience; representational content most directly reflects its intuitive conditions of satisfaction.

Because of this, the two-stage view yields natural answers to the objections to the Fregean view that were grounded in phenomenological adequacy. On the relationality objection: the two-stage view accommodates relationality by noting that there are certain specific and determinate properties—the perfect color properties—that are presented in virtue of the phenomenology of color experience. When Jack and Jill both have phenomenally green experiences in different environments, the two experiences have a common Edenic content, and so both are presented with perfect greenness. This captures the intuitive sense in which objects look to be the same to both Jack and Jill; at the same time, the level of ordinary Fregean and Russellian content captures the intuitive sense in which objects look to be different to both Jack and Jill. By acknowledging Edenic phenomenal content in addition to Fregean phenomenal content, we capture the sense in which perceptual phenomenology seems to be Russellian and relational.

On the simplicity objection: in the two-stage view, the simplicity of phenomenological structure is directly mirrored at the level of Edenic content. In Edenic content, there need be

no reference to normal causes, and no reference to experiences: instead, simple properties are attributed directly. The residual question for the Fregean view concerned how a complex Fregean content might be grounded in simple phenomenology. The two-stage view begins to answer this question. A given experience is most directly associated with a simple Edenic content, and this Edenic content is then associated with a Fregean content by the invocation of the matching relation. There is still an explanatory question about just where the matching relation comes from, and how it might be grounded: I address this question later in the chapter. But the two-stage view already gives us a skeleton around which we can build an explanatory connection between phenomenology and Fregean content.

On the internal unity objection: the two-stage view can accommodate the internal unity between visual and tactile experience of space by holding that the Edenic content of both visual and tactile experiences involve the attribution of perfect spatial properties (although the other perfect properties attributed by the experiences may differ). If so, then internal unity is present at the level of Edenic content. Further, the Fregean content of each will invoke the properties that match perfect spatial properties (in effect, the common typing of visual and spatial experiences is induced by the commonality in their Edenic content), and this common Fregean content will entail a common ordinary Russellian content. So the unity at the level of Edenic content will lead to unity at the level of ordinary content. Something similar applies to the case of representing the same color under different illumination; I will discuss this case in some detail shortly.

The two-stage view respects the insights of both the primitivist and the Fregean views in obvious ways. Like the original Fregean view, it can also respect certain key elements of dispositionalist and physicalist views. On the two-stage view, dispositions to cause relevant sorts of experiences still play a key role, not as the properties that are represented by experiences, but as a sort of reference-fixer for those properties. And the properties that are represented by the experience (at the standard of imperfect veridicality) are themselves plausibly physical properties, at least in the actual world. We might say that the view generates a broadly dispositionalist ordinary Fregean content and a broadly physicalist ordinary Russellian content.

9 Eden and Edenic content

What constraints are imposed by Edenic content?

The view I have proposed raises many questions. In the remainder of this chapter I address some of these questions, and in doing so flesh out a number of aspects of the view.

These include questions about Eden and Edenic content; about colors and color constancy; about matching and Fregean content; and about generalizing the model beyond the case of color. The order of these topics is arbitrary to some extent, so it is possible to skip to the topics that seem the most pressing.

A world with respect to which our visual experience is perfectly veridical is an *Edenic world*. (I defer until below the question of whether Edenic worlds are metaphysically possible.) It is natural to ask: what is the character of an Edenic world? A full answer to this question depends on a full analysis of the phenomenology of visual experience, which cannot be given here. But we can say a few things. As before, I will concentrate mostly on the aspects of phenomenology and representation associated with color, and will leave other aspects until later.

For any given experience, there will be many worlds with respect to which it is perfectly veridical. A visual experience—even a total visual experience corresponding to an entire visual field—typically makes quite limited claims on the world, and is neutral about the rest. For example, a visual experience typically presents things as being a certain way in a certain location, and is neutral about how things are outside that location. So, strictly speaking, in order to make an experience perfectly veridical, a world need merely be Edenic in certain relevant respects in a certain relatively limited area, and may be quite non-Edenic outside that area. Correspondingly, there will be a very large range of worlds that satisfy the relevant Edenic content. But here we can focus on what is *required* in order that the content be satisfied.

In a world that satisfies a typical Edenic content, primitive color properties such as perfect redness and perfect blueness are instantiated. Most often, visual phenomenology presents color as instantiated on the surface of objects, so an Edenic world will contain objects in which perfect colors are instantiated at certain locations on their surfaces. Strictly speaking, it will contain objects with certain perfect location-color properties: properties of having certain perfect colors at certain locations. Occasionally we have the phenomenology of volumes of color: as with certain transparent colored objects, for example, or perhaps with smoke and flames. In these cases, the corresponding Edenic world will have objects in which the relevant perfect colors are instantiated at locations throughout the relevant volume. It may be that sometimes we have the phenomenology of color not associated with objects at all: perhaps our experience of the sky is like this, just representing blueness at a certain distance

in front of us. If so, then a corresponding Edenic world will simply have perfect color qualities instantiated (by the world?) at relevant locations.⁹

From the fable at the beginning of the chapter, one might infer that Edenic worlds must meet a number of further constraints: perceivers must be directly acquainted with objects and properties in those worlds, illusion must be impossible, and there must be no microphysical structure. On my view this is not quite right, however. Edenic content puts relatively simple constraints on the world, involving the instantiation of perfect properties by objects in the environment, and these further constraints are not part of Edenic content itself. Their relation to Edenic content is somewhat more subtle than this.

Perfect color properties are plausibly *intrinsic* color properties. By virtue of presenting an object as having a perfect color at a certain location, an experience does not seem to make claims about how things are outside that location. So, when an object is perfectly red in Eden, it is this way by virtue of its intrinsic nature. In particular, it seems that an object can be perfectly red without anyone experiencing the object as perfectly red. The phenomenology of color does not seem to be the phenomenology of properties that require a perceiver in order to be instantiated. (The phenomenology of pain is arguably different in this respect, as I will discuss later.) It seems coherent to suppose that there is a world in which perfect colors are instantiated, but in which there are no perceivers at all.

One *could* hold a view on which, for an experience to be perfectly veridical, a subject must perceive the relevant perfect colors. On such a view, the character of visual experience is such that in addition to representing the presence of colors, visual experiences also represent the *perception* of colors. If one held this view, one would hold that no such experience is perfectly veridical unless the relevant perfect colors are perceived by a subject (the subject at the center of the relevant centered world), perhaps by direct acquaintance.

I am inclined to think that the character of visual experience is not like this, however. The phenomenology of color vision clearly makes claims about objects in the world, but it does not obviously make claims about ourselves and our perceptual relation to these objects. As theorists who introspect and reflect on how our phenomenology seems, we can say that on reflection it seems to us (introspectively) as if we are acquainted with objects and properties in the world. But it is not obvious that perceptual phenomenology itself makes such a claim:

⁹ Clark (2000) suggests that visual experience always involves the mere attribution of colors to locations rather than to objects. I find this suggestion phenomenologically implausible, but if it is correct, one could accommodate it by saying that Edenic worlds involve the instantiation of perfect color qualities by locations (or the instantiation of perfect color-location properties by the world), without requiring any special relationship between these qualities and objects.

to suggest that it does is arguably to overintellectualize perceptual experience. If perceptual experience does not make such claims, then the Edenic content of a visual experience will require that the relevant perfect properties are instantiated, but they will not require that we stand in any particular perceptual relation to those properties.

If this is correct, then in order to satisfy the Edenic content of an experience, a world must be Edenic in that perfect properties are instantiated within it, but it need not be a world in which we have not yet eaten from the Tree of Illusion. If an experience does not represent itself, it does not represent that it is nonillusory.

Likewise, a world that satisfies the Edenic content of an experience need not be one in which we have not yet eaten from the Tree of Science. The phenomenology of vision is arguably quite neutral on whether the world has the relevant scientific structure, as long as it also has primitive properties, and there is no obvious reason why a possible world could not have both.

To reinforce this view, we can note that the argument from the existence of illusions and of scientific structure to the nonexistence of perfect colors in our world was not a deductive argument. Rather, it was a sort of abductive argument: it undercut our reasons for accepting (instantiated) perfect colors, by suggesting that they are not needed to explain our visual experience. It remained *coherent* to suppose that primitive properties are instantiated in our world, but there was now good reason to reject the hypothesis as unnecessarily complex. On this view, eating from the Trees (by discovering the existence of illusions and scientific structure) did not directly contradict the Edenic contents of our experience, but it gave us good reason to believe that our world is not an Edenic world.

A more complete account of the Edenic content of color experience would require careful attention to all sorts of phenomenological details that I have largely ignored so far, such as the phenomenal representation of the distribution of colors in space, the fineness of grain of color representation, the different levels of detail of color experience in the foreground and background of a visual field, and so on. I cannot attend to, all of this here, but as a case study, I will shortly pay attention to one such detail, the phenomenon of color constancy.

What is the character of Edenic perception?

Even if perceivers are not presented in the Edenic content of an experience, it is natural to speculate about how perception might work in an Edenic world. One way to put this is to ask: what sort of world maximally reflects how things seem to us *both perceptually and introspectively*? Even if perception makes no claims about our perceptual experiences and our

perceptual relation to the world, introspection does. It seems to us, introspectively and perceptually, as if we stand in certain sorts of relation to the world. For this seeming to be maximally veridical, an Edenic world must contain subjects who stand in certain intimate relations to perfect properties in the world. We can call a world in which these seemings are maximally veridical a *pure Edenic world*.

Of course there are (possibly impure) Edenic worlds in which subjects perceive perfect colors via a mediated causal mechanism, at least to the extent that we perceive imperfect colors via such a mechanism in our world. But it is natural to think that this is not the best that they could do. It seems reasonable to hold that in Eden, subjects could have a sort of direct acquaintance with perfect colors. Perfect colors seem to be the sort of properties that are particularly apt for direct acquaintance, after all. And phenomenologically, there is something to be said for the claim that we seem to perceive colors directly. Certainly there does not seem to be a mediating causal mechanism, and one could suggest more strongly that at least introspectively, there seems not to be a mediating causal mechanism.

It is natural to suggest that in the purest Edenic worlds, subjects do not perceive instances of perfect color by virtue of having color experiences that are distinct from but related to those instances. That would seem to require a contingent mediating connection. Instead, Edenic subjects perceive instances of perfect colors by standing in a direct perceptual relation to them: perhaps the relation of acquaintance. Edenic subjects still have color experiences: there is something it is like to be them. But their color experiences have their phenomenal character precisely by virtue of the perfect colors that the subject is acquainted with. It is natural to say that the experiences themselves are constituted by a direct perceptual relation to the relevant instances of perfect color in the environment.¹⁰ We might say: in Eden, if not in our world, perceptual experience extends outside the head.

In the purest Edenic worlds, there are no illusions (if we take both introspection and perception to be maximally veridical, we conclude that things are just as they seem). In such a world, all color experience involves direct acquaintance with instances of perfect color in the environment. As soon as we eat from the Tree of Illusion, we have good reason to believe that

¹⁰ Is Edenic perception causal? Given that a perceptual experience consists of a relation of acquaintance with a perfect color property, is its character causally related to the perfect color property? This depends on subtle questions about the causation of relational properties by their relata. Compare: when a boy's first sibling is born, does this sibling cause the boy to be a brother? I am inclined to say yes, and to say the same thing about Edenic perception, holding that it involves a sort of unmediated causal relation. One could also say no, saying that this is a constitutive relation that is stronger than any causal relation. But even if so, there will at least be a counterfactual dependence of perceptual experience on perfect color properties in the world, by virtue of the constitutive relation.

we are not in such a world. But this need not cast us out of Eden entirely. There are somewhat less pure Edenic worlds in which there are illusions and hallucinations: perceivers sometimes have experiences as of perfect redness when the perceived object is perfectly blue, or when there is no object to be perceived. In these cases, the color experience cannot consist of a direct perceptual relation to an instance of perfect redness, as the subject stands in no such relation. Instead, it seems that the character of the experience is constituted independently of the properties of the perceived object.

In these impure Edenic worlds, an illusory or hallucinatory color experience involves a relevant relation to the property of perfect redness, without this relation being mediated by a relation to an instance of this property. (Something like this view is suggested as an account of hallucination in the actual world by Johnston 2004.) If so, then in such a world there may be phenomenally identical experiences (say, veridical and falsidical phenomenally red experiences) whose underlying metaphysical nature is quite distinct: one is constituted by a perceptual relation to a property instance in the subject's environment, and one is not. This picture is reminiscent of that held by some disjunctivists about perceptual experience in our world. We might say that in Eden, if not in our world, a disjunctive view of the metaphysics of perceptual experience is correct.

Is Eden a possible world?

Eden does not exist, but could it have existed? That is, is there a possible world in which there are perfect colors? Could God, if he had so chosen, have created such a world?

I am not certain of the answer to this question. But I am inclined to say yes: there are Edenic possible worlds.

To start with, it seems that perceptual experience gives us some sort of grip on what it would be for an object to be perfectly red, or perfectly blue. It would have to be exactly like *that*, precisely as that object is presented to us as being in experience. It seems that we can use this grip to form concepts of qualities such as perfect redness and perfect blueness (I have been deploying these concepts throughout this chapter). And there is no obvious incoherence in the idea that an object could be perfectly red, or perfectly blue. On the face of it we can conceive of such an object. So there is a *prima facie* case for believing that such an object is possible.

One can also reason as follows. There are good reasons to think that perfect redness is not instantiated in our world. But these reasons are empirical reasons, not a priori reasons. It was eating from the Tree of Illusion and the Tree of Science that led us to doubt that we live in an Edenic world. And eating from these Trees was an empirical process, based on

empirical discoveries about the world. Before eating from these Trees, there was no special reason to doubt that our experience was perfectly veridical. In particular, it is hard to see how one could be led to the conclusion that perfect redness is not instantiated by a priori reasoning alone (although see below). So the hypothesis that our world is Edenic seems at least to be conceivable, and it is reasonable to suggest that it cannot be ruled out a priori.

I have argued elsewhere (Chalmers 2002b) that this sort of conceivability is a good guide to metaphysical possibility. In particular, there is good reason to believe that if a hypothesis is ideally negatively conceivable, in that it cannot be ruled out by idealized a priori reasoning, then there is a metaphysically possible world that verifies the hypothesis. And there is even better reason to believe that if a hypothesis is ideally positively conceivable, in that one can imagine a situation in which the hypothesis actually obtains (in a way that holds up on idealized a priori reflection), then there is a metaphysically possible world that verifies the hypothesis.

The hypothesis that our world is Edenic (that is, that perfect colors are instantiated in our world) seems to be at least *prima facie* negatively conceivable (it cannot easily be ruled out a priori) and *prima facie* positively conceivable (we can imagine that it actually obtains). Furthermore, it is not clear how this hypothesis could be undercut by further a priori reasoning. If it cannot, then the hypothesis is ideally (negatively and positively) conceivable. If so, and if the thesis above is correct, then there is a metaphysically possible world that verifies the hypothesis. Verification is a technical notion from two-dimensional semantics (verification goes with primary intensions, satisfaction with secondary intensions), but the technicalities do not matter too much in this case (the primary and secondary intensions of perfect color concepts are plausibly identical, so that if a world verifies the hypothesis that perfect colors are instantiated it also satisfies the hypothesis). So, if this reasoning is correct, one can simply say: it is metaphysically possible that perfect colors are instantiated.

One could resist the conclusion either by denying that the Edenic hypothesis is conceivable in the relevant senses, or by denying the connection between conceivability in the relevant senses and possibility. Speaking for myself, I am reasonably confident about the latter, but I am not certain about the former. I do not see any obvious way of ruling out the Edenic hypothesis a priori, but I cannot be sure that there is no such way. We will see later that in the case of perfect pains, discussed below, there is arguably such a way. These considerations do not generalize to colors, but they make salient the possibility that other considerations might. For now, I am inclined to think that an Edenic world is metaphysically possible, but I am not certain of this.

Is there a property of perfect redness?

If what I have said so far is right, there is no *instantiated* property of perfect redness, but it is natural to hold that perfect redness may be an uninstantiated property. It seems that we have a grip on such a property in experience: we grasp what it would be for an object to have the property of perfect redness. Certainly, if an Edenic world is metaphysically possible, then objects in those worlds will be perfectly red, and it seems reasonable to conclude that they have the property of perfect redness. And even if an Edenic world is metaphysically impossible, one might still hold that there is such a property: it is just a necessarily uninstantiated property (like the property of being a round square). These issues will interact with one's views on the metaphysics of properties to some extent: for example, if one thinks that properties are just sets of possible objects, or if one thinks that properties are very sparse relative to predicates, one might resist some of the reasoning here. But overall I think there is a good *prima facie* case for thinking that there is a property of perfect redness.

If there is no such property, or if there is no metaphysically possible Edenic world, then some of the details in this chapter might have to change. If there is no metaphysically possible Edenic world, one cannot model the conditions of satisfaction association with perfect veridicality using sets of (or functions over) metaphysically possible worlds. And if there is no property of perfect redness, one cannot say that there is a content that attributes this property to an object. But even if so, one could understand the contents in other terms. For example, one could understand Edenic contents in terms of sets of epistemically possible scenarios rather than metaphysically possible worlds. Or one could understand Edenic conditions of satisfaction using something like Fregean concepts rather than properties. One could also regard Eden as some sort of mere world-model, not yet a possible world, but one that plays a key role in determining the ordinary Fregean contents of perception, via the requirement that the actual world must match the world-model in various respects. In this fashion numerous key elements of the two-stage model of perceptual content could be preserved.

If there is a property of perfect redness, what sort of property is it? It is most natural to conceive of perfect redness as a sort of simple, irreducible quality, one that might be instantiated on the surface of objects in some possible world. Perfect color properties might not all be maximally simple. For example, they might be seen as a sort of composition from simpler perfect properties, such as certain perfect unique hues (so that a particular shade of perfect orange may be a composite of perfect redness and perfect yellowness to certain degrees, and a certain amount of perfect brightness). But the underlying properties are naturally held to be irreducible.

In particular, it is natural to hold that perfect colors are not reducible to physical properties. If one accepts the earlier arguments that perfect color properties are not instantiated in our world, this consequence follows naturally. But even if one thought that perfect color properties are instantiated in our world, one could still argue that they are irreducible to physical properties, by analogs of familiar arguments concerning phenomenal properties.¹¹ For example, one could argue that one can conceive of a physically identical world in which they are not instantiated, and infer that such a world is metaphysically possible. Or one could argue that someone without color vision could know all about the physical properties of objects without knowing about their perfect colors.

Still, it is at least coherent to hold a view on which experiences have Edenic content that represents the instantiation of perfect color properties, and to hold that as a matter of empirical fact, perfect color properties are identical to certain physical properties (such as surface reflectances). On this view, our *concepts* of perfect color properties may be simple and irreducible concepts, but they pick out the same properties as those picked out by certain physical properties. Such a view would be analogous to certain “type-B” materialist views about phenomenal properties, according to which phenomenal properties are empirically identical to certain physical properties, because simple phenomenal concepts pick out the same properties as certain physical concepts. On the resulting view, experiences could be seen to have a Russellian phenomenal content that represents the instantiation of certain physical properties (although the experience does not represent these properties *as* physical properties). On this sort of view, our experiences might be perfectly veridical even in a purely physical world. I do not find this view plausible myself: it is vulnerable to the usual objections to Russellian physicalist views based on inversion scenarios, for example (requiring either strong externalism about phenomenology or arbitrary asymmetries among inverted communities), and it is also subject to the conceivability arguments above. But I think that there is at least an interesting variety of Russellian physicalism regarding phenomenal content in the vicinity.¹²

One could likewise hold a view on which perfect color properties are empirically identical to certain dispositional properties; or one could hold a view on which perfect color properties are distinct from physical and dispositional properties, but on which they

¹¹ Analog arguments of this sort are discussed in detail by Byrne (forthcoming). Byrne conceives these arguments as arguments for the irreducibility of color properties. I think the arguments work best as arguments for the irreducibility of perfect color properties.

¹² The version of Russellian physicalism about phenomenal content advocated by Byrne and Hilbert (2003) may be particularly close to this view.

metaphysically supervene on such properties.¹³ These views will be confronted with familiar problems: for example, the question of how to individuate the properties while still retaining plausible results about veridicality and illusion (for the view on which perfect colors are identical to or supervene on dispositional properties), and the questions of inversion and conceivability (for the view on which perfect colors supervene on intrinsic physical properties). But again, views of this sort are at least worth close attention.

Finally, it is possible to hold that perfect color properties are identical to certain mental properties, such as properties instantiated by one's visual field. This view agrees with the ordinary Edenic view that perfect colors are not instantiated by ordinary external objects, but holds that they are instantiated by certain mental objects (though they need not be represented *as* mental properties). The resulting view, a version of projectivism, does not suffer from the problems for the physicalist and dispositionalist views outlined above.¹⁴ I am inclined to reject this view myself, because of familiar problems with holding that mental objects instantiate color properties or their analogs (Chisholm's (1942) "speckled hen" problem, for example), and because the view becomes particularly hard to accept when extending beyond the case of color (it is hard to accept that mental objects instantiate perfect height, for example, of the sort that we represent in spatial experience). But the question of whether perfect properties might be instantiated in mental objects is at least well worth considering, and the corresponding version of projectivism might be able to accommodate many of the features of the two-stage view that I have been advocating.

For the remainder of this chapter, I will assume that perfect color properties are irreducible properties that are not instantiated in our world. But at least some aspects of the discussion may generalize to the other views I have outlined.

How can we represent perfect redness?

If perfect redness is never instantiated in our world, then we have never had contact with any instances of it. If so, one might wonder: how can perfect redness be represented in the content of our experiences?

Construed as an objection, this point turns on the tacit premise that representing a property requires contact with instances of it. In reply, one can note that we can certainly represent other uninstantiated properties (the property of being phlogistonated, Hume's

¹³ The views of Campbell (1993) and McGinn (1996) are at least closely related to the views on which perfect color properties supervene on intrinsic physical properties (for Campbell) and on dispositional properties (for McGinn).

¹⁴ The projectivist view of color defended by Boghossian and Velleman (1989) seems to be compatible with an Edenic view on which the perfect color properties are instantiated by a visual field.

missing shade of blue), and can even represent uninstantiable properties (being a round square). An opponent might suggest that these are complex properties whose representation derives from the representation of simpler properties, and so might suggest the modified premise that representing a *simple* property requires contact with instances of it. It is far from clear why we should accept this, however. For example, there seem to be perfectly coherent Humean views of causation on which we represent the simple property (or relation) of causation in our experience and in our thought, but in which there is no causation present in the world.

Certainly, there are cases in which representing a property crucially depends on contact with instances of it. But there many cases of representation that do not work like this. One can plausibly represent the property of being a philosopher without being acquainted with any philosophers. The same goes for causation, on the Humean view above. One might divide representations into those that are subject to Twin Earth thought experiments (so that twins in a different environment would represent different properties), and those that are not. Representations in the first class (including especially the representation of natural kinds such as water) may have content that depends on instantiation of the relevant property in the environment. But representations in the second class (including perhaps representations of philosophers and causation, at least if this representation does not involve deference to a surrounding linguistic community) do not depend on instances of the property in this way. In these cases, representation of a property comes not from instances of that property in the environment, but rather from some sort of internal grasp of what it would take for something to instantiate the property. It is plausible that representation of perfect redness falls into this second class.

Of course to say this much is just to respond to the objection, and not to fully answer the question. The residual question concerns just *how* our mental states get to have a given Edenic content. I will not try to answer this question here. We do not yet have a good theory of how our mental states represent any properties at all, and the cases of “narrow” representation, such as the representation of philosophers and causation above, are particularly ill-understood. To properly answer these questions, and the analogous question about Edenic content, requires a theory of the roots of intentionality.

I would speculate, however, that the roots of Edenic content lie deep in the heart of phenomenology itself. Horgan and Tienson (2002) have suggested that there is a distinctive sort of “phenomenal intentionality” that is grounded in phenomenology, rather than being grounded in extrinsic causal connections. It is not unreasonable to suppose that Edenic content is a basic sort of phenomenal intentionality—perhaps even the most basic sort. This

could be combined with a variety of views about the metaphysics of phenomenal intentionality. For example, one could hold that such intentionality is grounded in the projection of properties of certain mental objects, as on the projectivist view above. Or one could hold that the representation of Edenic content is even more primitive than this. If one is inclined to think that there is something irreducible about phenomenology, one might naturally hold that perceptual phenomenology simply consists in certain primitive relations to certain primitive properties: the presentation of perfect redness, for example. In any case, it is likely that understanding the roots of Edenic content will be closely tied to understanding the metaphysics of phenomenology.

10 Colors and Color Constancy

What about color constancy?

Color constancy is the phenomenon wherein instances of the same color in the environment, when illuminated by quite different sorts of lighting so that they reflect different sorts of light, nevertheless seem to have the same color. A paradigmatic example is a shadow: when we see a surface that is partly in shadow, although there is something different about the appearance of the shadowed portion of the surface, it often does not seem to us as if the object has a different color in the shadowed portion. One might say: although there is a sense in which the shadowed and unshadowed portions look different, there is also a sense in which they look the same. Certainly, the shadowed and unshadowed portions produce phenomenally distinct experiences, but we often do not judge that the object has a different color in those areas.

To say this much is to stay neutral on the representational content of the relevant experience. But it is natural to wonder just how the content of such experiences should be analyzed. In particular, it is natural to wonder how the two-stage model can handle such contents. To address this question, one can ask as before: how would the world have to be, in order for experiences of this sort to be perfectly veridical? A definite answer to this question requires a close phenomenological analysis. I will not give a full analysis, but I will outline some options.

It is useful to focus on the case of shadows. As an example, we can take a white floor on which an object casts a crisp dark shadow. I will take it that there are visual cues indicating that a shadow is being cast, so that we judge that the floor is still white at the relevant point, though we also judge that it is in shadow. What is the content of this experience? How would the world have to be, in order for the experience to be perfectly veridical?

The answer depends on how we analyze the phenomenology of the experience. To start, one might take either a *simple* or a *complex* view of the phenomenology. On the simple view, the apparent sameness in color between the shadowed and unshadowed area is not present in visual phenomenology at all. Rather, the sameness is detected only at the level of visual judgment, or perhaps at the level of other perceptual mechanisms whose contents are not reflected in phenomenology. For simplicity, let us say it is at the level of visual judgment. On this view, the phenomenal character of the experience of the floor may be the same as the phenomenal character of a floor where the relevant portion of the floor is painted the relevant shade of gray, and in which the floor is under constant illumination; it might also be the same as in a case where the floor is in shadow in the relevant portion, but where there are no cues. (We can stipulate that the last two cases involve exactly the same retinal stimulation, so that there is not much doubt that the resulting experiences are phenomenally identical.) On the simple view, the original shadow case will differ merely in that relevant cues lead to a judgment of sameness in that case but not in the others. The simple view will say something similar about all cases of color constancy: the constancy is present at the level of judgment, not at the level of perceptual experience.

The simple view is naturally associated with a view on which the local phenomenology of color experience is three-dimensional: the relevant experiences can be arranged in a three-dimensional color solid that exhausts the relevant dimensions of variation. Or at least it will hold that if there are further dimensions of variation, then variations due to shadows, illumination, and so on are not among them. On this view, the local phenomenology of perceiving the shadow will be the same as the local phenomenology of veridically perceiving an unshadowed object that is a relevant shade of gray. It is natural to hold that the Edenic content of such an experience involves the attribution of perfect greyness. So, on this view, the perfect veridicality of a shadow experience will require the instantiation of the relevant shade of perfect grayness in the object of perception. If we accept the simple view, then if a shadow is cast in a pure Edenic world (one without illusion), the color of the object will change.

On the simple view, what are the imperfect veridicality conditions of such an experience? An experience of the shadow will be correct iff the floor instantiates a property that matches perfect grayness. A property matches perfect grayness, to a first approximation, if it normally causes phenomenally gray experiences. If we take it that there is a canonical normal condition that involves unshadowed light, then this property will be something like a certain specific surface reflectance that the shadowed area of the floor does not instantiate, so the experience will be (imperfectly) falsidical. If we allow that there is a wide range of normal conditions

that includes both shadowed and unshadowed light, things are more complicated. I will discuss this complication further in the next section.

One other position compatible with the simple view holds that while the *local* phenomenology of seeing the partially shadowed floor is the same as the *local* phenomenology of seeing a partially grey floor without cues, the *global* phenomenology of the two cases is different (because of the difference in cues), and this difference in global phenomenology makes for a difference in conditions of veridicality. This view requires a certain anti-atomism about perceptual content: the veridicality conditions of an experience of a color at a location are not determined just by the local phenomenology associated with the location, but by the phenomenology of the entire visual experience. That is: two experiences can have the same local phenomenology but different local content, due to different global phenomenology. This view leads to a complicated further range of options about perceptual content, on some of which the shadow experience may end up being (imperfectly) veridical. These options end up roughly mirroring the options for the complex view that follows (which also postulates differences in local content, this time associated with differences in local phenomenology), so I will not discuss it further.

The alternative to the simple view is the *complex* view, on which the apparent sameness in color between the shadowed and unshadowed areas is present in some fashion in the visual phenomenology of seeing the floor. On this view, the experience of seeing the partially shadowed floor is phenomenally different from the experience of seeing a partially gray floor under uniform lighting, and the phenomenal difference is present in the visual phenomenology associated with the floor itself (and not merely in the experience of background cues). On this view, the presence or absence of cues makes a difference to the visual experience of the floor itself: one might say that the cues play a pre-experiential role and not just a pre-judgmental role.

This view is naturally associated with a view on which the local phenomenology of color experience is more than three-dimensional. For the sameness is accommodated in visual phenomenology; it is natural to hold that the color contents associated with the shadowed and unshadowed areas are in some respect the same. If local phenomenology were three-dimensional, and if differences in local content go along with differences in local phenomenology (the alternative that rejects this second thesis collapses into the anti-atomistic version of the simple view above), then this sameness in local content entails that the local phenomenology of seeing the shadowed and unshadowed white regions is exactly the same. That claim is not phenomenologically plausible. So the complex view suggests that the local

phenomenology of seeing color has more than three relevant dimensions of variation, with correspondingly more dimensions of variation in representational content.

On this view, the shadowed and unshadowed area will be represented as being the same in some respect: intuitively, both will be represented as white. They will also be represented as being different in some respect: intuitively, one will be represented as being in shadow and one will not. And these respects of sameness and difference will both be present in the phenomenology. One can argue that this view is more phenomenologically attractive than the simple view, in allowing phenomenological and representational differences between seeing something as shadowed white and as unshadowed white, on the one hand, and between seeing something as shadowed white and as unshadowed gray, on the other. I am inclined to favor the complex view over the simple view for this reason, although I think that the correct characterization of the phenomenology is far from obvious and neither view is obviously correct or incorrect.

If the complex view is correct, what should we say about the Edenic content of an experience of shadowed white? Phenomenologically, such an experience seems to characterize the intrinsic properties of a surface: if one takes the experience completely at face value, there seems to be an intrinsic (although perhaps temporary) difference between the shadowed and unshadowed parts of the floor. So it is natural to say that the Edenic content of the experience attributes a complex intrinsic property to the floor. One might see this property as the conjunction of two intrinsic properties: roughly, perfect whiteness and perfect shadow. That is, the Edenic content presents the floor as being perfectly white, infused in the relevant areas with a perfect shadow. This conjunctive treatment of perfect shadowed white is not mandatory: one could see the property as a certain mode of perfect white, rather than as a conjunction of perfect white with an independent perfect shadow property. But the conjunctive proposal has a certain phenomenological plausibility, insofar as one can see differently colored areas as subject to the same sort of shadow.

On this view, perfect shadows are things that can come and go in Eden, while the perfect color of an object stays the same. When a perfect shadow is cast on a perfectly white object, the shadow is *on* the object in the sense that it affects the intrinsic nature of the object's surface. Of course there are different sorts of shadows, corresponding to different degrees of shadowing, each of which can come and go while an object's perfect color stays the same. Strictly speaking, it is best to talk of shadow properties instantiated at locations on objects, rather than talking of shadows: while we sometimes have the phenomenology of seeing shadows as objects, it is arguable that more often we do not.

One might worry that this view cannot adequately capture the dimension of sameness between shadowed white and unshadowed gray. There is a clear respect in which these experiences are phenomenally similar, and one might argue that this respect corresponds to a representational similarity: perhaps one could say that the objects of such experience seem the same with respect to superficial color, or something along those lines. The representational claim is not obviously mandatory here, but if one accepts it, one might elaborate the Edenic model by saying that there is a respect in which any objects with perfect shadowed white and perfect unshadowed gray are similar to each other. One might say that both of these perfect properties entail perfect superficial grayness, for example. This might either be seen as a composite property, or simply as corresponding to another way of carving up the underlying multidimensional space.

What are the imperfect veridicality conditions of such an experience? Presumably an experience as of shadowed white is veridical iff its object has a property that matches perfect shadowed white; or, on the conjunctive treatment of shadowed white, iff it has a property that matches perfect white and a property that matches perfect shadow. The former is plausibly a physical property such as a certain surface reflectance (although see below). As for the latter, it will be a property that normally causes experiences as of the appropriate sort of shadow. It seems that no intrinsic property of surfaces is a good candidate here. Rather, the reasonable candidates are all relational: for example, the property of being subject to the occlusion of a light source to a relevant degree in the relevant area. This is a relational property rather than an intrinsic property, so it does not match the property of perfect shadow as well as it could. But with no intrinsic property being even a candidate, it seems that this property may match well enough. If so, then we can say the experience is imperfectly veridical iff the object has the relevant physical property (imperfect whiteness) and the relevant relational property (imperfect shadow). If it has one but not the other, one can say that the experience is imperfectly veridical in one respect but not the other.

One can extend something like this treatment to other cases of color constancy, and to cases of variation in illumination in general. One might hold that whenever there are relevant cues about illumination, these make a difference to the complex phenomenology of an experience with a corresponding difference in content. If the perceptual system is doing its job, then the object will be represented as having the same color, but it will also be represented as being different in some relevant respect, analogous to the presence or absence of shadows earlier. The difference in phenomenology seems to involve a difference in intrinsic (if temporary) properties, so the associated Edenic properties are intrinsic: one might call them perfect illumination properties (with the recognition that perfect illumination is

intrinsic rather than extrinsic). There will plausibly be a complex space of such perfect illumination properties, perhaps a three-dimensional space, and a corresponding space of matching imperfect properties (which may once again be relational properties, such as the property of being illuminated by certain sorts of light. Once we consider color and illumination together, we will plausibly have at least a six-dimensional space of complex Edenic properties in the vicinity, and a corresponding space of imperfect physical/relational properties.

One might wonder about the experience of darkness. What happens in Eden, if darkness falls? I am inclined to say that darkness is in some respects like the experience of shadow, but more all-pervasive. In particular, as darkness falls, darkness seems to pervade the environment, present at every location. The whole space appears to become dark. Objects do not seem to change their colors, exactly, although the representation of their colors may become much less specific, and it eventually becomes absent altogether (as does the representation of objects, in pitch blackness). So it is natural to say: in Eden, when things become dark, perfect darkness is present throughout the relevant volume of space, intrinsically altering that volume, although it need not alter objects' colors. In Eden, when darkness falls, perfect darkness pervades.

What are imperfect colors?

The imperfect colors are the properties that match the perfect colors (in our world), and whose instantiation or non-instantiation makes our color experiences veridical or falsidical. Just which properties are these? So far, I have said that these are the intrinsic physical properties that serve as the normal cause of experiences with the corresponding phenomenal properties. A first approximation suggests that these may be certain surface reflectance properties, or, better, the categorical bases of the relevant surface reflectance dispositions. But there are some tricky issues.

One tricky issue, stressed by Hardin (1987), arises from the fact that there is no such thing as a canonical normal condition for the perception of colors. Instead there is a wide range of normal conditions, including bright sunlight, muted cloudy light, shaded light, and so on. For a given subject, the same object may cause experiences with different phenomenal characters in each of these conditions. So it is not obvious that there will be any specific physical property that can be singled out as the "normal cause" of a given phenomenal character property.¹⁵

¹⁵ Hardin (1987) also stresses variations between normal perceivers, as well as variations in normal conditions. Variations between normal perceivers are no problem for the two-stage view, as matching is always relative to a

How we handle this issue depends on whether we take the simple view or the complex view of color constancy. On the complex view, as long as the mechanisms of color constancy work reasonably well, then the same object may cause experiences that are the same in certain key respects, while differing in other key respects. For example, a white object will cause an experience of shadowed white in shadowed conditions, and an experience of unshadowed white in unshadowed conditions. On the complex view in the previous section, the Edenic contents of these experiences attribute the same perfect color property (perfect whiteness) but different perfect illumination properties (perfect shadow and perfect unshadow). We can put this by saying that the experiences have the same phenomenal color property, and different phenomenal illumination properties. On this view, while a given object may trigger experiences with different phenomenal character in different conditions, these experiences will usually attribute the same phenomenal color (though different phenomenal illuminations), associated with the same perfect color property. So on this view, the wide range of normal conditions is not incompatible with the existence of a reasonably specific property that typically causes experiences with the relevant phenomenal color (that is, experiences that attribute the relevant perfect color) across the range of normal conditions.

If we take the simple view of color constancy, the issue is more difficult. On this view, a white object may cause quite different experiences under bright and shadowed light: let us call them phenomenally white and phenomenally gray experiences. On this view, there is no relevant phenomenal property that is shared by such experiences: any sameness in content enters only at the level of judgment. On this view, a phenomenally gray experience may be caused by a white object in one condition and by a gray object in another condition, where both conditions are equally normal. So it appears that on the simple view, there is no fine-grained intrinsic property that can serve as “the normal cause” of a phenomenally gray experience. (A similar issue could arise on the complex view, if it turns out that the mechanisms of color constancy are sufficiently unreliable.) And appealing to dispositional properties will not help, as the fine-grained dispositional properties of a white and a gray object differ as much as their intrinsic properties.

Here there are a number of possible reactions. One could hold that one condition (e.g. bright midday sunlight) is singled out as normal. One could hold that the matching imperfect property is not an intrinsic or dispositional property of the object but a (transient) relational property, such as the property of (currently) causing phenomenally gray experiences, or the

subject, and it is the normal cause for a given subject that determines the Russellian content of a color experience. At worst, this sort of variation has the consequence that the Russellian content of phenomenally identical color experiences in different subjects may represent different physical properties.

property of reflecting a certain sort of light, or the disjunction of being white under shadowed light, gray under unshadowed light, and so on. Or one could hold that it turns out empirically that no imperfect property matches perfect grayness, so that the (imperfect) Fregean content of such an experience determines no nontrivial Russellian content in the actual world (it is akin in certain respects to a sentence containing an empty description).

In my view the most plausible line for a proponent of the simple view to take is to hold that the normal cause of phenomenally gray experiences is a disjunctive or coarse-grained intrinsic property: one whose instances include white objects, gray objects, and any objects that cause phenomenally gray objects in some normal condition. On this view, a phenomenally gray experience of any such object will be veridical. This view has the advantage of capturing our intuitions that no such experience in reasonably normal conditions should be privileged over others, and that at least some of these experiences are veridical. The disadvantage of this view is that it suggests that the (imperfect) color properties attributed by color experiences are less fine-grained than one might have thought, and that in particular, a phenomenally gray and a phenomenally white experience do not attribute incompatible (imperfect) properties, even when they occur simultaneously. But on reflection, the consequence does not seem too bad: the incompatibility is still captured at the level of Edenic content, and if one takes the simple view and thinks of shadowed cases, it is reasonably intuitive that phenomenally gray and phenomenally white experiences might be compatible (for example, that both might veridically represent a white floor).¹⁶

Properties such as imperfect redness will be disjunctive in other respects. Color experience is most often caused by the reflection of light from objects, but it is also caused by the radiation of light from light sources, by the transmission of light from semi-transparent sources, and so on. The relevant cause of phenomenal color experiences in the first case will be something like a surface reflectance (or its categorical basis), but in other cases it will be something like a radiation profile (or its categorical basis). It seems reasonable to hold that color experiences of radiating objects and the like can be just as (imperfectly) veridical as those of reflecting objects. So imperfect redness is best seen as a disjunction of a range of

¹⁶ It might be objected that on this view, a phenomenally gray and white “striped” experience of an unstriped floor will be counterintuitively classified as veridical. This could be handled by saying that such an experience represents that the relevant areas are *differently* colored, a sort of relational color property. One could then say that the imperfect relation that best matches the relevant Edenic relation here is that of having different fine-grained surface reflectances, or something along those lines. Then the ordinary content of the experience will require that different fine-grained intrinsic properties are present in the relevant areas, but it will not take a precise stand on which fine-grained properties these are.

reflectance properties, radiation properties, and other properties that can serve as the relevant basis.

What are colors?

What does this view say about the nature of colors? Philosophers argue about whether colors (such as redness) are best seen as physical properties, dispositional properties, mental properties, primitive properties, or something else. So far I have taken no stand on this matter. What view of colors does the two-stage view suggest?

It is reasonable to hold that much of the issue here is terminological. We can acknowledge a role for properties of each of these sorts. Once we understand the precise role that each plays, we understand the substantive issues in the vicinity, whichever of them we choose to call “color”. That being said, the terminological issue is not wholly without content. There are certain core roles that we expect colors to play, and different properties are differently suited for the label “color” to the extent that they play more or fewer of these core roles.

On the two-stage view, the natural candidates to be called “colors” are perfect colors and imperfect colors. Both of these can be seen as playing one crucial role associated with colors: they are properties attributed in color experiences. Perfect colors are attributed in Edenic contents, and imperfect colors are attributed in ordinary contents. Perfect colors play certain further core roles that imperfect colors do not: we seem to be acquainted with their intrinsic nature in color experience, and the perfect colors arguably stand in relevant intrinsic structural relations to each other in a way that imperfect colors do not.

Still, perhaps the core role of colors is that they are the properties whose instantiation is relevant to the truth of ordinary color attributions. That is, an utterance of “that apple is red” will be true if and only if the apple instantiates redness. Furthermore, it is natural to hold that some apples really are red. The two-stage view is partly driven by the thesis that some ordinary color experiences are veridical (even if they are not perfectly veridical). It seems equally reasonable to hold that apples really are red (even if they are not perfectly red). If so, this suggests that redness is not perfect redness but imperfect redness.

So I am inclined to say that color terms, in their ordinary uses, designate imperfect color properties. Just which properties these are depends on how matching is understood, and to the extent that matching is somewhat indeterminate, the designation of color terms may be somewhat indeterminate. But I am inclined to think that our ordinary uses of color terms designate certain disjunctive physical properties, with properties such as surface reflectance properties among the disjuncts. The physical properties designated by ordinary color terms

will be relatively coarse-grained, but there will be more fine-grained physical properties in the vicinity, which we might regard as the different shades of these colors.¹⁷

Of course, one can also reasonably use color terms to refer to Edenic properties (“perfect redness”), phenomenal properties (“phenomenal redness”), and maybe to other properties as well. On this view, there are multiply interlocked families of properties: the perfect colors, the imperfect colors, possibly further families of imperfect colors associated with different notions of matching, and the phenomenal colors. As long as we understand the complex relationships between these families, and the roles that each can play, not much of real substance rests on the question which of these families is the true family of colors.

There is nothing especially original or distinctive about the view of the ontology of color that emerges from the two-stage view. In identifying colors (in the core sense) with physical properties, the resulting ontology of colors may be very similar to that of the physicalist about color. And the various families of color properties that are introduced may also be acknowledged by the primitivist. Although primitivists about color identify perfect color properties with the colors, they may also recognize that physical and dispositional properties play some of the roles of the colors. For example, Maund (1995) says that terms such as “red” refer in their core sense to the perfect colors (which he calls “virtual colors”), but also refer in an extended or metonymic sense to the physical properties that I have called imperfect colors. So the ontology recognized by this sort of primitivist view is not dissimilar to that recognized by the two-stage view.

What is distinctive about the two-stage view is not its associated ontology of colors but rather, its view of perceptual content. On the primitivist view, experiences have a single content (an Edenic content) that determines their veridicality. On the two-stage view, experiences have two layers of content, an Edenic content that reflects their phenomenology and a Fregean content that determines their veridicality. It is this two-layered view of content that is responsible for most of the explanatory power of the two-stage view.

¹⁷ If an advocate of the simple view of color constancy takes the line I suggested, on which experiences with specific phenomenal shades attribute relatively coarse-grained physical properties, then we will have three relevant sorts of physical properties. There will be highly coarse-grained properties (with some indeterminacy at the edges) that are the referents of terms such as “red”; there will be somewhat coarse-grained properties that are attributed by specific color experiences; and there will be fine-grained physical properties of which these coarse-grained properties can be seen as ranges or disjunctions. Probably the fine-grained properties are the best candidates to be called the “shades”. The main costs are that we lose a tight correspondence between physical shades and phenomenal shades, and that there will not turn out to be a specific physical shade that qualifies as “unique red” (though there may still be unique phenomenal red, and unique perfect red). If we take the complex view of color constancy (and if the mechanisms of color constancy are sufficiently reliable), then these problems are avoided.

Is this indirect realism?

One might worry that this view is a form of indirect realism about color perception. According to standard indirect realism, we perceive objects in the world only indirectly, by virtue of directly perceiving certain intermediate objects such as sense-data, which opponents see as a “veil of perception” that cuts off perceivers from the external world. The two-stage view I have outlined is certainly not a variety of standard indirect realism, as it does not invoke any intermediate objects as objects of perception. But one might worry that it is a form of indirect realism about the perception of *properties*. In particular, one might suggest that this is a view on which instantiated color properties (that is, imperfect color properties) are perceived only indirectly, by virtue of directly perceiving perfect color properties.

This objection invokes the relation of *perception* between subjects and properties. This relation is analogous to the relation of perception between subjects and objects: it is natural to say that when I veridically perceive a green square in the environment, I both perceive the square and perceive its greenness. So far in this chapter, I have focused on the relation of perceptual representation, but not on the relation of perception. These seem to be different relations: one can perceptually represent an object or a property without perceiving it (in a hallucination, for example).

The standard view of the perceptual relation between subjects and objects holds that it is a *causal* relation: to perceive an object is roughly to have a perceptual experience that is appropriately caused by the object (and perhaps that has a phenomenal character that is appropriately related to the character of the object). The standard view of the perceptual relation between subjects and properties is presumably something similar: to perceive a property is roughly to have a perceptual experience whose phenomenal character is appropriately causally related to an instance of that property (and perhaps whose phenomenal character represents the instantiation of the property, or otherwise “matches” the property in some fashion).

If we adopt this standard view of the perceptual relation, there is no threat of indirect realism. In a typical veridical experience of a green object, the phenomenal character of my experience is causally related to the relevant instance of physical greenness, and represents the instantiation of physical greenness in its Russellian content. By contrast, the phenomenal character of my experience is not causally related to any instance of perfect greenness, as there are no such instances. So it seems that on the two-stage view, as much as on other views of perceptual experience, we perceive imperfect colors directly, and not by virtue of perceiving any other property.

It is true that on the two-stage view, perception is not as “direct” as perception could be. There is a sense in which perception in Eden is more direct than it is in our non-Edenic world. In Eden, perception works by direct acquaintance, and there need be no mediation between objects and properties perceived and a perceptual experience. In our world, there is complex causal mediation. This does not entail that our perception is *perceptually* mediated, though, as on the indirect realist view.

We might say that in Eden, an especially strong perceptual relation obtains, one that we might call *perfect perception*. Perfect perception of an object or property requires unmediated acquaintance with the object or the property, and perhaps also requires that the object or the property is itself a constituent of one’s perceptual experience. By contrast, *imperfect perception* requires only the appropriate sort of causal connection to an object or a property. As before, it is plausible to suggest that if we took the deliverances of both perception and introspection fully at face value, we would conclude that we live in an Edenic world in which we perfectly perceive objects and properties in that world. But after the fall from Eden, there is no perfect perception; there is just imperfect perception.

We might call this view not *indirect realism* but *imperfect realism*. Our acquaintance with the world is not as direct as it would be in Eden, and perception does not reveal the intrinsic nature of things in the way that it does in Eden. But this is so for any causal theory of perception. Perception on the view I have outlined is no more and no less imperfect than on most causal theories. The idea of Eden just brings out the contrast, for all these theories, with the kind of perfect perception that we cannot have in our non-Edenic world. One might yearn for the kind of perfect contact with the world that we had in Eden, but after the fall, we have learned to live with the imperfection of perception.

11 Matching and Fregean Content

What is matching?

The notion of matching serves as the bridge between Edenic content and ordinary content. An experience is imperfectly veridical when its object has properties that match the perfect properties attributed by the experience. But what is it for a property to match a perfect property? To a first approximation, we can say that a property matches a given perfect property (for a given subject) if the property is the normal cause of the associated phenomenal property (in that subject). But this is clearly just a first approximation.

A basic constraint is that at most one imperfect property can match a given perfect property. Or at least, at most one imperfect property can match a perfect property for a subject

at a time. Different imperfect properties can match the same perfect properties for different subjects, and probably for the same subject at widely separated times. (Strictly speaking, we should say that matching is a three- or four-place relation involving subjects and times, but I will usually leave the subject and the time in the background.) But we need at most one matching property for a subject at a time, in order that the ordinary Russellian content of an experience can attribute a property to its object. Of course it could be, for all we have said, that matching imperfect properties are often disjunctive properties, or determinable properties with many different determinates. And it may be that sometimes there is no imperfect property that matches a given perfect property.

Matching is best understood as a holistic relation. Rather than saying that imperfect redness is the property that normally causes phenomenal redness, one can say that the set of imperfect color properties is that three-dimensional manifold of properties that serves as the normal causal basis for the associated three-dimensional manifold of phenomenal color properties. This requires that there is a mapping from imperfect properties to phenomenal properties such that in many or most cases a given imperfect property will normally cause the associated phenomenal property, but this relation need not hold in all cases. If there are exceptions associated with certain imperfect properties in the manifold (such as Kripke's "killer yellow", a shade of yellow that always kills the perceiver if observed; or perhaps a Humean missing shade of blue that is never instantiated in our world for a lawful reason), this will not stop the manifold as a whole from matching, and the imperfect property will still be associated with a corresponding phenomenal property. When this mapping associates an imperfect property with a phenomenal property that attributes a given perfect property, we can say that the imperfect property matches the perfect property.

It will be clear that the notion of matching is a vague and messy one. One source of messiness arises from the issue discussed above: there is no precise delineation of the class of normal conditions. Even if there were such a delineation, there is no precise criterion for when a property causes an experience often enough in these conditions to count as its normal cause. Further, there is more to matching than normally causing an associated phenomenal property. We have seen that there are structural constraints, such as the constraint that imperfect color properties fall into the same sort of three-dimensional manifold as perfect color properties. There are also categorical constraints, such as the constraint that imperfect color properties be intrinsic properties if possible. And it is presumably desirable that (imperfect) color properties be properties that can stand in the sort of relations to (imperfect) spatial properties that perfect color properties stand in to perfect spatial properties.

One could attempt to encapsulate all these constraints and others in a full and precise definition of matching, but I am not optimistic about the prospects for such a definition, any more than I am for definitions of other philosophically important notions such as perception and knowledge. An alternative approach is simply to say: matching is that relation M such that, necessarily, an experience is imperfectly veridical iff its objects have properties that bear M to the properties attributed by its Edenic content. In effect, this notion exploits our relatively pretheoretical grip on imperfect veridicality, along with an independently grounded notion of Edenic content (explained in terms of perfect veridicality, say), to explicate the notion of matching.

Of course this explication does not say anything substantive about what matching involves. For a substantive characterization, one has to rely on our judgments about the (imperfect) veridicality and falsidicality of experiences. We do have quite clear judgements in many cases. And it is plausible that we judge experiences to be veridical precisely when objects in the world instantiate certain properties, properties that correspond in some fashion to the perfect properties in the Edenic content of our experiences. Even if we cannot give a full account antecedently of what this correspondence consists of, there is good reason to believe that it is present, and one can say quite a lot about what it involves in specific cases, as we have done above. For example, it usually seems to require normal causation of an associated phenomenal property, and there are other constraints as suggested by various cases. As in the case of analyzing knowledge, there will probably be no straightforward articulation of necessary and sufficient constraints, but nevertheless the consideration of cases can help us to flesh out the constraints in the vicinity.

One might worry that this characterization taken together with the two-stage view will be circular. The two-stage view says that an experience is imperfectly veridical iff its objects have properties that match the relevant perfect properties. The characterization above says that matching is that relation M such that an experience is imperfectly veridical iff its objects have properties that bear M to the relevant perfect properties. There is no circularity, however. In the project of explication, we have a prior grip on the notion of imperfect veridicality, and we use this prior grip in order to explicate the notion of matching. Via this explication, we theoretically characterize a relation M . One can then use relation M for certain theoretical purposes, if one likes. At the very least, we can appeal to it in analyzing the relationship between imperfect and perfect conditions of veridicality. One might go further and hold that, metaphysically, for an experience to be imperfectly veridical is for its objects to bear M to the relevant perfect properties. Or one might hold that epistemically, our intuitive judgments about imperfect veridicality are mediated by a tacit prior grasp of M . I am cautious

about making such further claims here, although I think there is something to them. But in any case, there is no more circularity here than in any other case where one uses a pretheoretical notion to help characterize a theoretical notion, which one then may use to help give a theoretical account of the pretheoretical notion.

Of course our judgments about (imperfect) veridicality are not always clear. There are many cases in which we are not sure what to say, or in which we are tugged in two different directions. Sometimes these judgments are cleared up on a certain amount of rational reflection, but sometimes they are not. When they cannot be cleared up in this way, the natural thing to say is that the relevant case is a vague case of imperfect veridicality. The vagueness of imperfect veridicality will give rise to a corresponding vagueness of matching: it will be vague whether the object in question instantiates a property that matches the relevant perfect property. There may be different ways of precisifying the notion of imperfect veridicality, which will give rise to corresponding precisifications of the notion of matching. But some vagueness and messiness in the notion of matching is just what we should expect, given the vagueness and messiness of imperfect veridicality.

Is Fregean content phenomenologically adequate?

Although the two-stage view has a clearer grounding in phenomenological structure than the original Fregean view, one might still worry about its phenomenological adequacy. The Edenic content of an experience (in which the two-stage content is grounded) seems to nicely mirror the structure of the phenomenology. But the imperfect Fregean content does not. In particular, there is nothing discernible in the phenomenology of visual experience that obviously corresponds to matching. Certainly, it is hard to see that there is any clear phenomenology of “normal causation” in a typical visual experience. And to the extent that matching is messier and more complex than a notion based on normal causation, it seems all the more distant from the phenomenology. For example, we have seen that matching can often be vague, as can the associated Fregean content. But the phenomenology itself need not be vague; or, if it is vague in some respects, it need not be vague in relevant respects. For example, it is plausibly vague in some cases whether an object has a property that matches perfect redness. But the associated phenomenally red experience may be quite precise, with the phenomenology of precisely presenting a specific property of the object. So one may ask, as we did before, whether this Fregean view is phenomenologically adequate.

Here, I think, one should concede that matching does not correspond directly to any element of the visual phenomenology. The phenomenology of visual experience is the same in our world and in Eden. The presentation of an Edenic world does not (or need not) involve

attribution of normal causation and the like. So the phenomenology of ordinary visual experience does not (or need not), involve this either. Perhaps there are some experiences that present causal and dispositional relations, but it seems wrong to say that every ordinary color experience does this.

Where does matching come from, then? I think the answer is clear: it comes from the *inferential role* of visual experience. The content of a mental state need not be something that one can read off the intrinsic properties of its vehicles. There is good reason to believe that quite generally, mental content is tied to inferential role. This is especially so in the case of Fregean content, which was introduced by Frege to mirror the cognitive and inferential significance of thought and language. A belief that “Hesperus is Phosphorus” has a very different inferential role from a belief that “Hesperus is Hesperus”, and this difference in inferential role is reflected in a difference in their inferential content. It is even possible to define the Fregean conditions of satisfaction of a belief partly in terms of the belief’s inferential role, such as the conditions under which a subject will rationally accept or reject the belief, given information about the world (see Chalmers 2002a for such an account).

Beliefs are not the only mental states that have inferential roles. Perceptual experiences also have an inferential role, broadly understood. Just as one belief can serve as grounds for accepting or rejecting another belief, a perceptual experience can likewise serve as grounds for accepting or rejecting beliefs, and more generally for guiding our knowledge about the world. Most obviously, one can *endorse* a perceptual experience, yielding a perceptual belief about the character of one’s environment, and that belief can be used to accept or reject other beliefs in turn. For example, when one has a phenomenally red experience as of an object in one’s environment, this can be used as grounds for accepting a belief that there is a red object in front of one. One would not normally call this relation between experience and belief an “inference”, but it can be seen as a sort of quasi-inferential relation.

Just as with belief, the inferential role of a perceptual experience can be analyzed in part by asking: when given information about how things are in the world, will a subject accept or reject the perceptual experience? That is, will they accept or reject the belief that things are as they perceptually seem to be? If one takes an example, such as a subject having a phenomenally red experience as of an object in front of them, one finds a specific pattern of judgments. If the subject discovers that there is really no object in front of her, she will reject the experience: things are not as they seem. If she discovers that there is an object in front of her but it has the sort of physical make-up that usually causes phenomenally green experiences (only causing phenomenally red experiences this time due to unusual lighting), then she will reject the experience: again, things are not as they seem. But if she discovers

that the object in front of her has the sort of make-up that usually causes phenomenally red experiences, then she will accept the experience: at least in the relevant respect, things are as they seem.

In effect, the core inferential role of a perceptual experience is reflected in the pattern of judgments about veridicality and falsidicality that the subject of such an experience makes, or, more strictly, in the pattern of judgments that they should rationally make. And we have already seen that this pattern of judgments closely corresponds to the Fregean content above. The pattern of judgments does not require that objects in the environment have any specific property, such as a surface reflectance, or even perfect redness. It requires only that the property be the property that plays the appropriate causal role. So as in the case of beliefs, this Fregean content closely mirrors the experience's inferential role.¹⁸

Here, we can respond to the charge of phenomenological adequacy by rejecting the claim that phenomenal content must precisely mirror phenomenological structure. Phenomenal content can equally be grounded in inferential role.

Of course, a proponent of the original Fregean view could have made the same response (as I did in response to a similar worry in Chalmers 2004a). So how is the two-stage view any better in this respect? To see the difference, recall where things stood at the end of section 5. It was not clear that the objections from phenomenological adequacy had any knockdown force, but they raised the issue of a serious explanatory incompleteness in the Fregean view. Fregean content is supposed to be a sort of phenomenal content, such that necessarily, an experience with the same phenomenology has the same Fregean content. But the presentational phenomenology of visual experience does not simply wear its Fregean content on its sleeve. So there needs to be some explanatory story about how Fregean content is related to the phenomenology of the experience, and why it is that any experience with that phenomenology will have this Fregean content.

It is this explanatory story that the two-stage view provides. The presentational phenomenology of an experience immediately grounds an Edenic content. The Fregean content is grounded in the Edenic content by virtue of inferential role. The subject is immediately presented, in visual phenomenology, with an Edenic world. But a rational subject need not hold the world to an Edenic standard. In effect, a rational subject will use the

¹⁸ If someone is doubtful that experiences have Fregean content (perhaps holding that there is only Edenic phenomenal content and ordinary Russellian nonphenomenal content), it is this pattern of judgments about veridicality, and the corresponding inferential role, that gives the best reason to believe in it. There is no doubt that experiences are associated with this sort of pattern of judgments of veridicality, and there is no obstacle to our using this pattern to ground a notion of experiential content.

Edenic phenomenology of a phenomenally red experience to ground the claim that the object in front of them is *red*, but she need not make strong claims about the intrinsic nature of redness. That is left open: if the subject discovers that objects with property P typically cause red experiences, then she will decide that those objects are red, and that if the original object has property P, then the original experience was veridical. In effect, the presentational phenomenology of the experience serves as direct ground for the first stage of the two-stage view (the Edenic content), and as indirect ground for the second stage (matching the Edenic content) by virtue of inferential role.

Is Fregean content phenomenal content?

Once we observe that ordinary Fregean content derives from inferential role, this may raise another worry: is Fregean content really *phenomenal* content? The mere fact that Fregean content does not completely mirror phenomenological structure here is no objection, as the definition of phenomenal content does not require this sort of mirroring. However, the definition does require that any experience with the same phenomenology has the same phenomenal content. And one may worry: if inferential role is extrinsic to phenomenology, could not two phenomenally identical experiences have different inferential roles, yielding distinct Fregean contents?

Of course, there is an obvious sense in which phenomenally identical experiences can have different inferential roles. For example, if I believe that red snakes are poisonous and you do not, then relevantly similar visual experiences in the two of us might produce quite different beliefs. But this difference in inferential role need not be a difference in the core aspect of inferential role that is relevant to defining Fregean content. This core aspect involves the subject's pattern of judgments of veridicality and falsidicality associated with the experience. More precisely, it turns on whether the subject should rationally accept or reject the experience (that is, judge that things are or are not as they perceptually seem to be) when given relevantly complete information about the world. Two subjects may have the same pattern of judgments here despite different beliefs. For example, in the case above, both subjects may well have the same rational dispositions to accept or reject the experience, given full information.

Fregean content will be phenomenal content as long as the same experience rationalizes the same pattern of judgments, given relevant information, in all subjects. (Of course there may be differences in an associated *actual* pattern of judgments due to cognitive limitations, but a rational inferential role idealizes away from such limitations.) This will be the case as long as: (i) every phenomenally identical experience has the same Edenic content; (ii) every

subject should rationally accept an experience, given relevant information, iff (according to that information) the relevant object has properties that match the properties attributed in the Edenic content; and (iii) the matching relation is the same for all subjects.

I think there is good reason to accept (i) and (ii). We have already seen that Edenic content is a sort of phenomenal content. I think, further, that the match-involving inferential role is rational for any subject with a perceptual experience. Such an experience presents a world with a certain distribution of Edenic properties, and rational judgments of veridicality should go with whether objects in the world have properties that match those properties. What is not so clear is whether one should accept (iii). We have already seen that the notion of matching is somewhat vague and imprecise. Could there not be subjects whose equally rational judgments invoke somewhat different matching relations, perhaps held to somewhat different standards in each case?

For example, one might suggest that before the fall from Eden, the inferential role of our experiences required a strict standard of matching. Perhaps an Edenic subject would judge an experience falsidical if they discovered that its object merely has an imperfect property that serves as its normal cause. However, there is good reason to hold that even our Edenic counterparts have dispositions such that *if* they were to discover that their world is non-Edenic, they would still judge their experiences to be (imperfectly) veridical when their objects have the relevant imperfect properties. After all, when we discovered that our world was non-Edenic, these were the judgments that we made. So there is reason to believe that Fregean inferential role is present even in Eden.

Still, one can ask whether there *could* be a rational subject (whether in Eden or outside) who has such a strict standard of matching that they will accept a phenomenally red experience only if the relevant object is perfectly red? If such a subject discovers that the world is non-Edenic, they will reject all their color experiences as falsidical. For such a subject, the relevant standard of matching would be the strict standard of identity: a property matches perfect redness iff it is perfect redness. Certainly there could be a subject that has an actual pattern of judgments like this. The relevant question is, could this pattern of judgments be as rational as the pattern of judgments that we have been discussing? The answer is not obvious.

Likewise, we can imagine subjects who make different judgments in difficult cases. For example, let us assume the simple view of color constancy. Then one subject might judge a phenomenally red experience to be veridical iff the relevant object has the property that causes such experiences in bright sunlight. Another subject might judge such an experience to be veridical iff the relevant object has any property in the range that might cause the

experience in some normal condition. And another subject might judge that no such experience is veridical as there is no single specific property that plays the right role. The question then is: could these patterns of judgment be equally rational?

Finally, we can consider a possible difference between visual and olfactory experience. We do not usually judge olfactory experiences to be veridical or falsidical. We do not say that a rotten egg smell is veridical iff there is sulphur dioxide nearby and falsidical iff there is not, for example. This is not because the phenomenology of smell is not representational: intuitively, it seems to represent that certain smells are present in the world. It is just that we are not inclined to make judgments of veridicality and falsidicality; at best, we make judgments of misleadingness or otherwise. On the other hand, perhaps there could be subjects who make judgments of veridicality or falsidicality for phenomenally identical olfactory experiences. For example, one can imagine that if dogs could make judgments, this is what they would do! One could diagnose this by saying that for those subjects, but not for us, there are properties in the environment that match perfect smells. The question is: are both patterns of judgment equally rational?

It is possible to say no in all these cases. One might hold that one pattern of judgments in these cases is rational and that the others are not. For example, one could argue that in the first case there is some irrationality in holding the world to an Edenic standard, and that in the second case it is irrational to reject a color experience when its object has a property that normally causes that sort of experience.

One could also hold that in at least some of these cases, insofar as it is possible for corresponding experiences to rationalize different patterns of judgment, there will be a corresponding difference in the phenomenology. For example, in the third case above, one could suggest that the phenomenology of olfactory experience in dogs and humans differs: perhaps dogs have a more strongly presentational phenomenology of smell, for example. More generally, one might hold that certain differences in the character of presentational phenomenology might go along with differences in the associated standard of matching. In such cases, the existence of different rationalized patterns of judgment will be no obstacle to Fregean content serving as a sort of phenomenal content.

My own view is that it is not obvious that phenomenology underdetermines the standard of matching, but it is not obvious that it does not. Whether it does or not depends on difficult questions about the rational role of perception, and also about its presentational phenomenology, that I cannot adjudicate here. But I think that it is at least a live possibility that the standard of matching is underdetermined, and that there could be distinct equally rational patterns of judgment associated with the same sort of experience in different subjects.

If this is so, then what follows? One could say that the phenomenally identical experiences have distinct Fregean contents (in which case Fregean content is not phenomenal content), or one could say that they have the same highly indeterminate Fregean content (in which case imperfect veridicality is highly indeterminate). But I think that the best thing to say in this event is that these experiences have the same *unsaturated* Fregean content. This content is one that is satisfied iff the relevant object has properties that match the relevant primitive properties. However, the standard of matching is left unspecified by this unsaturated content, so the condition of satisfaction is in a certain sense incomplete. To yield a complete condition of satisfaction, the unsaturated content needs to be saturated by specifying a standard of matching. The resulting saturated Fregean content will yield a reasonably determinate condition of imperfect veridicality.

On this view, only unsaturated Fregean content, and not saturated Fregean content, will be phenomenal content.¹⁹ This is a step back from the original view of Fregean content as phenomenal content, as an unsaturated Fregean content is not a complete condition of satisfaction. That is, it is not the sort of thing that is true or false absolutely in a scenario. Correspondingly, the unsaturated Fregean content of an experience does not determine whether or not the experience is imperfectly veridical in its environment. What determines imperfect veridicality is a saturated Fregean content, which is not fully determined by phenomenal character.

What determines saturated Fregean content, if not phenomenology? One natural answer is inferential role, here conceived as something that might vary independently of phenomenology. In the different subjects above, phenomenally identical experiences play different inferential roles, yielding different saturated Fregean contents. In effect, the different inferential roles in different subjects (as reflected in a pattern of veridicality judgments) determine different standards of matching. In this way phenomenology and inferential role together determine a saturated Fregean content, and a condition of imperfect veridicality.

An alternative suggestion is that saturated Fregean content is determined not by inferential role but by a *standard of assessment* that is extrinsic to the subject. On this view,

¹⁹ Gideon Rosen suggested that for our counterparts on Psychedelic Earth, where experiences drift in a way that is completely unrelated to the environment (and where subjects know this), the experiences will have no inferential role at all, and so will not even have unsaturated Fregean content. However, it seems that if these subjects were to discover that the environment contains drifting properties that match the drifting experiences, they would then judge their experiences to be veridical. (It is precisely because they have discovered that there is no such match with the environment that they reject their experiences as a guide to the external world.) So it seems that these subjects still have the inferential disposition to endorse their experiences if they discover that the matching relation obtains. If so, this suggests that their experiences have Fregean content.

in effect, one could evaluate the same experience as either veridical or falsidical at different standards of assessment. We have already introduced dual standards of perfect and imperfect veridicality; on this view, there will be a range of different standards in the vicinity of imperfect veridicality. This range of standards will correspond to a range of different standards of matching. To evaluate an experience with an unsaturated Fregean content, we must tacitly introduce a standard of matching. This standard will determine a saturated Fregean content, and according to this standard the experience will qualify as veridical or falsidical.

One can then say: our ordinary notion of veridicality tacitly invokes a certain standard of matching, one that is reasonably although not completely determinate. With this standard fixed, phenomenally identical experiences will have the same saturated Fregean contents. However, there might have been different evaluators with a slightly different notion of veridicality, corresponding to a different standard of matching. With that standard fixed, phenomenally identical experiences will also have the same saturated Fregean contents. But these contents will differ from those associated with our standard. One might say that on this view, any given experience is associated with a whole range of (saturated) Fregean contents, depending on the corresponding notion of veridicality. Each of these Fregean contents could be seen as a sort of phenomenal content.

The two suggestions—according to which saturated content is determined by inferential role or by an external standard—yield somewhat different treatment of cases. Take a subject whose inferential role holds her experiences to the Edenic standard: upon discovering that the world is non-Edenic, she rationally rejects her perceptual experiences. On the former view, we will say that her experience is nonveridical: it is her own rational inferential role that determines ordinary veridicality. On the latter view, we will say that her experience is veridical: it is our standards that determine the veridicality of an experience (according to the meaning of our term “veridical”). However, if she or someone sharing a similar standard were to say that her experience is “nonveridical”, they would also be correct: they express a slightly different notion of satisfaction with their terms “veridical”. On reflection I find the second suggestion somewhat more plausible and intuitive than the first, although the matter is far from obvious.

In any case, whichever view we take, one can say the following. The phenomenal character of an experience determines an Edenic phenomenal content, and it determines an unsaturated Fregean phenomenal content. According to the unsaturated phenomenal content, an experience is veridical iff the relevant object has properties that match the relevant Edenic properties. Once combined with a standard of matching, this unsaturated content determines a

saturated Fregean content. This saturated Fregean content may or may not be phenomenal content, depending on what view one takes on the questions above. If one thinks that there is only one rational standard of matching associated with the phenomenal character of the experience, then the saturated Fregean content will be a phenomenal content. If one thinks that the associated standard of matching depends on a contingently associated inferential role in the subject, then the saturated Fregean content will not be a phenomenal content. If one thinks that the standard of matching is determined by an external standard of assessment, then the Fregean content will be a phenomenal content, but there will be a range of other Fregean contents associated with different standards of veridicality.

The choice between these three alternatives turns on difficult and subtle issues that I will not try to resolve here. But in any case, we can be confident that phenomenal character determines Edenic content and unsaturated Fregean content. The status of saturated Fregean content as phenomenal content remains an open question.

12 Beyond Color

I have concentrated on the content of color experience, but I think the two-stage model has much broader application. Here I will much more briefly discuss the extension to other aspects of perceptual experience.

Spatial experience

Apart from colors, the most salient properties attributed in visual experience are spatial properties. Does the two-stage model of phenomenal content generalize to these? I am inclined to think that it does.

One might think that spatial experience is more amenable to a straightforward Russellian treatment than color experience. But as Thompson (forthcoming) has argued, many of the same problems arise. A natural candidate for the Russellian content of spatial experiences involves the attribution of spatial properties such as that of being in a certain (absolute) location. But this content obviously cannot be phenomenal content, as a phenomenally identical experience could be had by a subject light-years away from that location. A natural next suggestion is a Russellian content involving the attribution of relative spatial properties (or relations, or relational property radicals) such as being six feet in front of the perceiver. But this cannot be phenomenal content either. In principle, a phenomenally identical experience could be had by a perceiver who is (and has always been) twice as big, in an environment where everything is twice as distant. Such an experience would not plausibly

attribute the same relative spatial property; it would more plausibly attribute the relative property of being 12 feet away.

One might then move to more relativized spatial properties, such as the property of being twice as distant or twice as big as some other object. Or one might suggest that phenomenal content can at least attribute shape properties, such as being square or circular. But as Thompson argues, similar problems arise. There could conceivably be an “El Greco” world in which everything is stretched 10 times in one direction compared with our world, but in which structure and dynamics are otherwise isomorphic. In such an environment, phenomenally square experiences would normally be caused by (what we call) long and thin rectangles. Further, there is good reason to think that such experiences would be veridical: certainly, if we found that we inhabit a corner of the universe that is locally stretched in this fashion relative to the rest of the universe, we would not conclude that our spatial experiences are falsidical. Rather, the natural thing to say is that phenomenally square experiences attribute different properties in these environments: (what we call) squareness in one environment, and a certain sort of rectangularity in another.

A more extreme case along these lines is given by a Matrix scenario, in which phenomenally identical subjects have been hooked up for their lifetime to a computer simulation of the world. I have argued elsewhere (Chalmers 2003b) that such subjects are not massively deluded about the world. Their beliefs such as “there are tables”, “I have hands”, and “that is square” are true; it is just that the underlying metaphysics of their environment is not what they expect (in effect, it is an underlying computational metaphysics). The same can be argued for their perceptual experiences: their experiences as of red square objects are as veridical as ours. However, such experiences need not be of (what we would call) square objects: there need be nothing square inside the computer. At best, there are objects with some very different property: we might call this “virtual squareness”. In any case, if phenomenally identical spatial experiences can be veridical in an environment that is spatially utterly unlike our own, this suggests that the phenomenal content of these experiences does not involve the attribution of ordinary spatial properties.

In this way, one can argue against Russellian views of spatial phenomenal content in ways that directly parallel our earlier arguments in the case of color experience. The natural alternative is a Fregean view of spatial phenomenal content. On this view, spatial experiences have Russellian content, attributing spatial properties and relations, but this content is not phenomenal content. Rather, phenomenal content involves a Fregean mode of presentation of spatial properties and relations: roughly, these are determined as that manifold of properties and relations that serves as the normal causal basis for the corresponding manifold of spatial

experiences. On this view, the Fregean content of a spatial experience is one that will be satisfied if the object has a property that normally causes the relevant sort of spatial experience (or if it has a complex of properties each of which normally causes the relevant sort of spatial aspect of the experience). One can then raise worries about the phenomenological adequacy of this view, motivating a two-stage view of spatial phenomenal content.

On the two-stage view, spatial experiences have an Edenic content that attributes perfect spatial properties: perfect squareness, perfect rectangularity, and so on. Arguably, even an Edenic content does not attribute absolute spatial properties, but just relative properties. It is not clear that we have the phenomenology of being presented with absolute spatial properties, and one can make a case that even in Eden, there could be phenomenally identical veridical experiences at different locations. But we do have the phenomenology of being presented at least with absolute shapes, and relative distances. So the Edenic contents of our experience will attribute perfect properties of this sort. It is plausible that these properties are not instantiated in our world (though arguing this takes a bit more work than in the case of color). If not, then our experiences are not perfectly veridical.

Our spatial experiences may nevertheless be imperfectly veridical, by virtue of their objects instantiating imperfect spatial properties: those that match perfect spatial properties. These will be the properties that serve as the normal causal basis for our spatial experiences. Imperfect veridicality will be associated with a corresponding ordinary Fregean content, one that is satisfied iff relevant objects have properties that match the relevant perfect spatial properties. Phenomenally identical experiences will have the same Edenic contents, and the same ordinary Fregean contents (setting aside issues about standards of matching), but may have different ordinary Russellian contents, because different properties may match the relevant perfect properties in different environments.

The Matrix provides a good illustration. The subjects here do not have perfectly veridical experiences, but they have imperfectly veridical experiences, by virtue of the fact that relevant matching properties (virtual squareness and the like) are instantiated in their environment. So subjects in the Matrix may share Edenic spatial contents with us, and may share ordinary Fregean contents also, but they will have different ordinary Russellian contents.

Of course the two-stage model of spatial experience needs to be elaborated in numerous respects to handle all sorts of aspects of spatial content: for example, perspective, angle, size constancy, mirror reflections, and the like. But there is reason to think it can help explain certain phenomena. For example, it is better suited than the original Fregean view to

accommodate internal connections between spatial representation in visual and tactile experience. On the original Fregean view, it might seem that there can be no internal connection, as the normal causes of visual spatial experience are not constrained to be the normal causes of tactile spatial experience. On the two-stage model, however, one can argue that the phenomenology entails that tactile and spatial experiences involve the attribution of common perfect spatial properties in their Edenic content. If so, then the matching imperfect properties will be constrained to be the same, thus grounding an internal connection between tactile and spatial experience.

It is a further question how this model should be extended to the representation of time and motion. I am inclined to say that the two-stage model can be extended to time as well as to space, though this turns on subtle issues about the metaphysics of time. A natural suggestion is that the Edenic content of temporal experience requires A-theoretic time, with some sort of true flow or passage. Our own universe may not instantiate these perfect temporal properties, but it may nevertheless instantiate matching B-theoretic properties (involving relative location in a four-dimensional “block universe”) that are sufficient to make our temporal experiences imperfectly veridical, if not perfectly veridical. The representation of motion could be treated in a similar way.

One might go so far as to suggest that Eden is a world with classical Euclidean space, and an independent dimension of time, in which there is true passage and true change. Our own world is non-Euclidean, with time and space interdependent, and with pale shadows of perfect passage and change. On this view, Einstein’s theory of spacetime was one more bite from the Tree of Science, and one more step in our fall from Eden.

The experience of objects

Our initial characterization of the Russellian contents of visual experience characterized them as having the following form: object O has color C at location L. In the case of color and location I have argued that this Russellian content is not phenomenal content, and have proposed a two-stage Fregean treatment instead. In the case of color, we have seen that the relevant Russellian content is also not plausibly phenomenal content. Does this mean that we should also give a two-stage Fregean treatment of the representation of objects?

A natural first suggestion is that experiences of objects have an Edenic content involving the representation of certain specific perfect objects: for example, perfect object O has perfect color C at perfect location L. However, this suggestion is implausible on reflection. In particular, it is implausible that the perfect veridicality of an experience of an object requires any particular Edenic object to be present. It seems that even in Eden, there could be two

phenomenally identical experiences of different objects. The phenomenology of object experience seems to present us directly with objects, but it does not seem to acquaint us with their intrinsic nature in a sense over and above acquainting us with their colors, shapes, and so on. If it did, then the phenomenology of object experience would be quite different from what it is: experiences of different tennis balls would typically have quite different phenomenal characters, for example. But the experience of objects does not seem to be this way.²⁰

Because of this, it is more natural to hold that even the Edenic content of object experience is existential. For example, one might hold that the Edenic content of an experience of a red sphere is satisfied iff there is a perfect sphere at the relevant location that is perfectly red. No specific object is required for the satisfaction of this content. On this view, Edenic content is not especially different from ordinary content in the representation of objects, so the two-stage model has no special role to play.

Still, there may be a further role for the two-stage model to play. One might hold that a merely existential characterization of phenomenal content does not fully respect the directness of an experience of an object.²¹ According to this objection, experience does not merely present that there *is* an object at a certain location with a certain color: it presents that *that* object is at a certain location with a certain color.

I think one might accommodate this suggestion without moving all the way to object-involving phenomenal contents, however. The phenomenology of perception does not seem to reveal the intrinsic haecceitistic natures of objects, but it does seem to present us with objects directly. To account for this, one can suggest that the experience of objects involves demonstrative modes of presentation.

In Eden, one is directly acquainted with objects, and no mediation is involved. One can simply demonstrate an object as *this* object, and acquaintance does the rest. This sort of reference is analogous to the unmediated way we refer to ourselves in our world, with “I”, or perhaps to the unmediated way in which we ostend our conscious experiences. An Edenic content might correspondingly have the form [that is C at L], where “that” is a primitive demonstrative, C is a perfect color, and L is a perfect location. The demonstrative here does

²⁰ As in note 1, a disjunctivist view about phenomenology may hold that the phenomenology of experiences of different objects differs in precisely this way. Such a disjunctivist view might hold that the Edenic content (and perhaps the non-Edenic phenomenal content) of an experience is object-involving. On my view, however, reflection on Eden suggests that a “naive realist” view of perception does not require disjunctivism about phenomenology. In Eden, a sort of naive realism about perception is correct, but this entails only disjunctivism about the metaphysics of experience (as discussed in section 9), not disjunctivism about phenomenology.

²¹ This sort of worry about the existential characterizations of perceptual content is canvassed by Campbell 2002 and Martin 2002. A demonstrative view of perceptual content is suggested by Burge 1991.

not build in the identity of the object, any more than the notion of “I” builds in a specific person: the same demonstrative could in principle refer to different objects, just as “I” can refer to different people. But neither is it associated with a substantive criterion of application. When the demonstrative has an object, it simply picks out the object directly, as *that* object. In the two-dimensional model, one could say: in Eden, one can refer directly to perceived objects as entities at the center of a centered world.

This Edenic content respects the direct presentational phenomenology of our experience of objects. But it is not clear that it has application outside Eden. In our world, we are not directly acquainted with objects outside ourselves: mediation is always involved. So our epistemic grip on objects is not as direct as it is in Eden, and the primitively demonstrative aspects of Edenic content are arguably not satisfied. Nevertheless, we stand in a weaker relevant relation to objects in our world: the relation of perception. One might say that by virtue of standing in this relation, the objectual aspects of our experience are *imperfectly satisfied*. There will be an associated condition of imperfect satisfaction. An object imperfectly satisfies the experience iff it is the object perceived with the experience: that is, if it is connected to the experience via an appropriate causal chain. One can think of this as a nonprimitive demonstrative condition of satisfaction: it comes with substantive requirements, but is grounded in a primitive connection to the experience itself. So, in effect, the objectual phenomenology of the experience can be perfectly or imperfectly satisfied: perfect satisfaction turns on primitive acquaintance, and imperfect satisfaction requires at least a mediated perceptual connection.

The imperfect satisfaction conditions of an experience can be seen as a sort of (ordinary) Fregean mode of presentation, picking out the object that the experience is appropriately connected to. The experience as a whole will be imperfectly veridical iff the object that is appropriately connected to the experience has properties that match the relevant perfect properties. On this view, the ordinary Fregean content of the experience will involve a connectedness condition of this sort, and it will determine in turn a Russellian content involving the relevant objects and its imperfect properties. Of course, the ordinary Fregean content is not a perfect mirror of the phenomenology: as usual, the phenomenology does not seem to involve reference to a causal condition, or reference to the subject’s experience. But this is just what we expect: Edenic contents mirror the phenomenology, and associated Fregean contents capture veridicality conditions after the fall.

If one takes this view, one will class so-called “veridical hallucinations” (hallucinations that happen to mirror the environment in front of one) as not really veridical at all. In these cases there is no object that one is perceiving, so the Fregean content is not satisfied and an

object-involving Russellian content is not determined. An alternative route to this result (Searle 1983; Siegel forthcoming) is to suggest that experiences have existential contents that attribute the relational property of being perceived with the relevant experience to the relevant object. Arguably, however, suggesting that this relational property is attributed along with color and location does not respect the subjunctive intuition that things could have been as they perceptually seem to be, even had there been no perceivers in the vicinity. By contrast, putting the perceptual requirement in the mode of presentation of the object allows this subjunctive intuition to be respected.

There is perhaps one other role for the two-stage model in the representation of objects. The phenomenology of vision seems to present a world that is carved into objects at its joints. One does not simply perceive a distribution of mass and color: one perceives objects on top of other objects, each of which may be articulated into objectual parts. Depending on one's metaphysical views, one may think that the world does not respect this articulation into objects. One might think that macroscopic objects do not exist in the world's basic ontology, or one might give their existence some highly deflationary treatment, on which their individuation is a matter of convention or conceptual scheme, or on which there is no deep fact of the matter about when there is an object or when there is not. But even if one's metaphysics is deflationary about objects, one's phenomenology is not. So perhaps, for our visual experiences to be perfectly veridical, there would have to be real, first-class, non-relative objects in the world. One might say that in Eden, there are perfect objects. If our world's ontology does not have perfect objects, or at least if it does not have perfect objects corresponding to the apparent objects of ordinary perception, then our experiences are not perfectly veridical in this respect. But they may nevertheless be imperfectly veridical, by virtue of there being appropriately arranged matter in the environment, or by virtue of the environment's satisfying some other deflationary condition. Once again, Eden sets the standard, and our imperfect world can only match it.

Other sensory modalities

The two-stage model can naturally be extended from visual experience to auditory and tactile experience. The details of these extensions depend on a careful analysis of the phenomenology of these experiences, combined with analysis of judgments about veridicality. But there is reason to believe that the model outlined in the case of vision will apply.

The phenomenology of auditory experience, at a first approximation, seems to represent certain sounds as being present at certain locations. For example, in a musical experience, the phenomenology might suggest that a sound with a certain pitch, timbre, volume, and so on is

being produced at a certain approximate location in front of me. As in the case of color, there are physical properties that one might plausibly identify with various pitch, timbre, and volume properties, and that one might hold to be attributed in an ordinary Russellian content. But these properties depend on the environment of the experience, and it seems that phenomenally identical experiences could have different Russellian contents of this sort. So one can move to a Fregean phenomenal content in these cases, and then, to respect phenomenological adequacy, hold that this content is grounded in the matching of an Edenic phenomenal content.

In Eden, one may hold, there are perfect sounds, with perfect middle-C pitch, perfect loudness, and so on. We grasp these simple intrinsic properties in our experience, but they are not instantiated in our world. Instead, in our world there are simply physical events such as air disturbances, with associated physical properties that match the Edenic properties. This is enough to make our auditory experiences imperfectly veridical, if not perfectly veridical.

Something similar goes for tactile experience. In Eden, objects may be perfectly smooth, or perfectly slimy, or perfectly velvety. These are intrinsic properties of objects or their surfaces, and we seem to be acquainted with these properties in our experience. But in our world there are just complex physical substitutes for these properties, such as imperfect sliminess and imperfect velvetiness. This is enough to satisfy the ordinary Fregean content of our tactile experiences, if not the Edenic content, and enough to make our tactile experiences imperfectly veridical.

Olfactory and gustatory experiences are trickier. The phenomenology of smell and taste seems to be representational. Intuitively, an olfactory experience represents that a certain smell is present in one's environment, perhaps in a certain broad location. A gustatory experience represents that something with a certain taste is in one's mouth or throat or on one's lips. But at the same time, we do not usually assess experiences of smell and taste for veridicality, and the notion of an illusory olfactory or gustatory experience does not get a strong grip on us. Certainly, there can be smell experiences that are caused by properties that do not normally produce such experiences, and the same applies to taste experiences (imagine a rewiring of the connection between receptors and the brain, for example). But it does not seem natural to describe such experiences as illusions. It is slightly more natural to speak of olfactory and gustatory hallucinations, when an experience is generated for reasons quite independent of external objects. But the intuition is not strong.

Taste and smell differ in this way from hearing and touch. We certainly assess auditory experiences for veridicality, and speak of auditory illusions if there is not a sound being produced where there seems to be. This way of speaking is less common in the case of touch,

as touch seems to be the most reliable of the sensory modalities, but we can nevertheless make good sense of the idea of a tactile illusion or hallucination. An object might feel smooth although it is not really smooth, or one might feel that an object is present when there is no object at all. In these cases we have no hesitation in classifying a tactile experience as falsidical. But in the case of taste and smell, one hesitates. I suspect that this is partly because we use taste and smell much less to gather information about our environment than we do hearing and touch, and partly (perhaps correspondingly) because the presentational element of their phenomenology is less striking.

Still, there is some presentational phenomenology in the experience of smell and taste. We seem to have some grip on intrinsic qualitative properties that are presented, although it is somewhat less obvious than in the case of vision that the phenomenology presents intrinsic properties of objects or of the environment as opposed to intrinsic qualities of experiences (or corresponding relational properties of objects and environment). Overall, though, I am inclined to say that olfactory and gustatory experiences have Edenic contents: the former presents perfect smells as being present in one's environment, and the latter presents perfect tastes as being instantiated in one's mouth.

It is the ordinary content of these experiences that is problematic. It is plausible that there are physical properties that normally cause the relevant olfactory and gustatory aspects of experiences, so one might think these would be the imperfect smells and tastes attributed in the ordinary content of these experiences. But because our assessments of veridicality are very unclear in these cases, it is likewise unclear whether these physical properties count as matching the relevant Edenic properties. In these cases, the standard for matching seems somewhat different from the case of vision and hearing, perhaps because of a difference in presentational phenomenology, or perhaps just because we apply a different standard because of different pragmatic purposes. So the status of ordinary Fregean and Russellian content in these cases is unclear. But we can nevertheless invoke Edenic content to help characterize the phenomenology.

Bodily sensations

What about bodily sensations, such as the pain experiences, itches, hunger, and orgasms? On the face of it, these have a strong presentational phenomenology. The experience of pain, for example, seems to present a certain painful quality as being instantiated in part of one's body, such as one's ankle. The experience of an itch seems to present a certain itchy quality as being presented on one's skin. In the phenomenology, these qualities seem to have a highly distinctive intrinsic qualitative nature. So it is natural to hold that bodily sensations have an

associated Edenic content, attributing Edenic properties such as perfect painfulness and perfect itchiness to locations in one's body.

There are two complications in this case. The first resembles the complication in the case of smell and taste. We do not generally assess bodily sensations for veridicality or falsidicality. Perhaps in an extreme case such as phantom limb pain, we are somewhat inclined to say there is some sort of falsidical pain hallucination. But we are not really inclined to speak of pain illusions, or of illusory itch experiences. If we did, we would probably be talking of a case where we mistake the phenomenal character of an experience, not where we mistake its object. As in the case of smell and taste, it seems that there are physical properties (such as tissue damage and the like) that normally cause the relevant experiences. But we are not especially inclined to say that when these properties are absent, an experience as of pain or as of an itch is falsidical. Even if there is no associated tissue damage, for example, we are not inclined to say that an intense pain experience is illusory. So the ordinary Fregean and Russellian content of these experiences seems somewhat unclear, in the same way as in the case of smell and taste.

A related complication concerns the Edenic content of bodily sensations. What are perfect pains like, in Eden? That is, what sort of properties need to be instantiated in one's body in order for a painful experience to be perfectly veridical? Here there are conflicting requirements. First, the properties seem to be intrinsic properties whose nature we grasp in experience. The phenomenology of pain in one's ankle seems to attribute a quality that is intrinsic to one's ankle. But second, the properties seem to have a strong connection to experience itself. Can one conceive of one's ankle being in perfect pain without anyone experiencing the pain? It is not clear that we can. In this respect the phenomenology of pain is quite different from the phenomenology of color, where we have no trouble conceiving of an object being perfectly colored even though no one ever experiences its color. But this seems strong to suggest that perfect pain is a relational property, as its instantiation places requirements on how things are outside the object in which it is instantiated.

Is the property of perfect pain intrinsic or relational? Neither answer is entirely comfortable. If perfect pain is an intrinsic property of an ankle, it seems that its instantiation should be independent of whether an experience is present. But it is not clear that unexperienced perfect pain is conceivable. But if perfect pain is a relational property, what relational property could it be: the property of causing a painful experience, or of having such-and-such intrinsic quality perceived in a painful experience? Neither of these seems apt to the phenomenology. Furthermore, the former seems to make too little a claim about what is going on in one's ankle, and the latter seems vulnerable to the objection that came up in the

intrinsic case: we do not seem to have a grip on any relevant intrinsic quality here that we can conceive instantiated in the absence of a painful experience.

Perhaps the best answer is the following: perfect pain is an intrinsic property, but one whose instantiation entails the existence of an associated painful experience, or of associated phenomenal pain. We might think of it as an intrinsic property that, if instantiated, necessarily “broadcasts” further constraints on the world. In effect, it is an intrinsic property that stands in a necessary connection to distinct intrinsic properties of experience. In effect, it is a property whose instantiation brings about necessary connections between distinct existences.

If this property could be instantiated, problems would follow. It is not clear that there can be necessary connections between distinct existences of this sort. It seems plausible that for any conceivable or possible situation in which an intrinsic property is instantiated in one’s ankle, it should be conceivable or possible that the property is instantiated in an arbitrarily different context. But it is not conceivable or possible that there is perfect pain without pain experience. The natural conclusion is that perfect pain cannot be instantiated: there is no possible world in which there is perfect pain, and on reflection it is not even conceivable that there is perfect pain. In effect, the instantiation of perfect pain places incoherent requirements on the world.

This does not entail that there is no property of perfect pain. There are other properties whose instantiation is impossible and inconceivable: that of being a round square, for example. One might hold that perfect pain is like this. On this view, one has a grip on the property of perfect pain, based on one’s experience. But one does not need to eat from the Tree of Illusion or the Tree of Science to know that perfect pain is not instantiated: one can know this simply on sufficient reflection. Perhaps there can be matching intrinsic properties (without the relational constraint), or matching relational properties (without the intrinsic constraint), but no property can play both roles. Still, one may hold that the property exists, and one can hold that it is attributed in the Edenic content of our pain experiences.²²

In effect, the Edenic content of pain sets a standard that is not just hard to meet, but impossible to meet. There are related instantiated properties, to be sure: that of causing painful experiences, for example, or having a certain sort of tissue damage. But because these fall so far short of playing the role of perfect pain (the former is not intrinsic, and the second has no strong connection to experience), one might suggest that they fail to match perfect pain. It is arguably because of this that we do not judge that the instantiation of these

²² Adam Pautz explores an idea like this in forthcoming work.

properties yields veridicality or falsidicality of pain experiences. The standard set by Eden is sufficiently high that there is little point holding the world to it.

What goes for pain goes also for other bodily sensations, such as the experience of itches, hunger, and orgasms. One finds the same combination in these cases: phenomenology seems to present an intrinsic property, but one that cannot be instantiated without a corresponding experience. The natural conclusion is that the perfect properties cannot be instantiated at all. One might suggest that this model applies in some other domains: for example, one might suggest that gustatory experiences present properties that cannot be instantiated except while being tasted, so to speak. If this were so, it could help to explain our reluctance to assess such experiences as veridical or falsidical. The phenomenology here is less clear than in the case of pain, and it is not obvious whether the claim of a necessary connection to experience is correct. But the analogy between the cases at least deserves attention.

It may be that some other Edenic properties that we have considered are not just uninstantiated but uninstantiable. For example, one might hold that perfect time (involving the flow of time, or a moving now) is incoherent, perhaps for McTaggartian reasons. Or if one is sufficiently deflationary about objects, one might hold that perfect objects cannot exist in any possible world. Nevertheless, the impossibility of satisfying these contents does not automatically stop them from acting as a regulative ideal. Here, the impossible might serve to regulate our experience of the actual.

High-level properties

One might try to extend this model beyond the representation of simple properties such as color and shape in experience, to the representation of high-level properties, such as that of being a duck, or being happy. It is plausible that representing such properties can make a difference to the phenomenology of experience (Siegel, forthcoming). It is not clear that the phenomenal content of this sort of experience is easily analyzed using the two-stage model. One difficulty is that the deployment of *concepts* often plays a key role in such experiences, and that the content of the experience is inherited from that of an associated concept rather than being determined by the two-stage model. When we see something as a book, or as a duck, for example, it is plausible that the associated phenomenal content is inherited from the content of our concept of a book or of a duck. And in these cases, we do not seem to have any grip on distinct perfect and imperfect veridicality conditions.

Still, there are a few cases where the two-stage model is at least tempting. For example, there is a phenomenology of moral experience, and it is arguable that moral properties such as being good or bad can be represented in perception. One might naturally suggest that for

moral experiences to be perfectly veridical, relevant objects would have to have perfect moral properties: the sort that are objective, intrinsically motivating, and so on. But it is arguable that in our world (and perhaps in every possible world), no such properties are instantiated. If so, our moral experiences cannot be perfectly veridical. But there are various properties (including response-dependent properties, community-relative properties, and so on) that arguably match these properties well enough. If so, our moral experiences can be imperfectly veridical. There has been no perfect goodness since the fall from Eden, but we can at least be consoled by imperfect goodness in the world.

13 Conclusion

On the view I have presented, the most fundamental content of perceptual experience is its Edenic content. Other aspects of content such as ordinary Fregean and Russellian content can be seen as deriving from Edenic content, with the aid of the matching relation and the contribution of the environment. To understand the role of perceptual experience in representing the world, one needs to understand all these levels of content. But to understand the phenomenology of perceptual experience in its own right, understanding Edenic content is the key.

We have seen that the Edenic approach yields a very useful tool in doing phenomenology. To characterize the phenomenology of an experience, it is often helpful to characterize the sort of world in which that experience would be perfectly veridical. To do this, one sketches relevant aspects of the character of Eden. Doing this does not eliminate the need for thorough phenomenological investigation, and it does not solve the many associated hard methodological problems, but it at least provides an analytic tool that gives us some purchase in characterizing the contents of consciousness.

I am inclined to think that Edenic content may also give us an entry point for understanding the metaphysics of experience. I have said little in this chapter about how it is possible for experiences to have Edenic contents, or about which of Edenic content or phenomenal character is the more fundamental. My suspicion is that neither is more fundamental than the other. It may be that perceptual experience is fundamentally equivalent to the presentation of an Edenic world. If so, then if we can understand how the presentation of an Edenic world is possible, then we will understand perceptual phenomenology.

Appendix: The Two-Dimensional Analysis of Perceptual Content

The Fregean and Russellian content of perceptual experience can be modeled using the two-dimensional framework for the analysis of content. This analysis is not required in order to make use of the notions of Fregean and Russellian content (which can be understood intuitively as in the text), but it helps in order to make the use of these notions more precise, and in order to analyze certain subtleties that arise. It also can help us to shed light on the relationship between the content of perception and of belief.

In the linguistic version of the two-dimensional framework, expression tokens are associated with two intensions, or functions from possible worlds to extensions. The Fregean content of an expression is associated with a primary intension: a function from centered worlds to extensions, where a centered world is a world marked with a designated individual and a designated time (intuitively, these represent the perspective of the subject who utters the expression). The Russellian content of an expression is associated with a secondary intension: a function from (uncentered) possible worlds to extensions.

For example, the primary intension associated with “Hesperus” might be an intension that in a given centered world, picks out the bright object that has been visible in the appropriate location in the evening sky around the center of that world. The secondary intension associated with “Hesperus” might be a function that picks out Venus in all worlds. Likewise, the primary intension associated with “Hesperus is Phosphorus” might be an intension that is true in a centered world iff the bright objects that have been visible at certain positions in the evening and morning skies around the center of that world are identical. The secondary intension associated with “Hesperus is Phosphorus” might be an intension that is true in a world iff Venus is Venus in that world.

This framework can be extended to the contents of mental states such as beliefs (Chalmers 2002a), and the discussion above suggests that it can also be extended to the contents of perceptual experiences. The Fregean content of a perceptual experience is associated with a primary intension. For example, the condition on extension associated with phenomenal redness can be modeled as a function from centered worlds to properties: in a given centered world, it picks out the property that normally causes phenomenally red experiences in the subject at the center of that world. The condition of satisfaction that is associated with a specific experience of a colored object can be represented as a function from centered worlds to truth-values: in a given centered world, it will be true iff there is a relevant object in the environment of the individual at the center that instantiates, at the relevant

location relative to that individual, the property that normally causes phenomenally red experiences in that individual.

Likewise, the Russellian content of an experience is associated with a secondary intension. For example, the Russellian content associated with phenomenal redness can be seen as an intension that picks out physical redness in all worlds. And the Russellian condition of satisfaction that is associated with a specific phenomenally red experience of an object *O* might be seen as an intension that is true at a world iff object *O* is physically red in that world.

Why do I say that Fregean and Russellian contents are associated with primary and secondary intensions, rather than that they are primary and secondary intensions? This is because I want to leave open the possibility that these contents have a more fine-grained structure than functions from worlds to extensions. For example, it is reasonable to hold that the Russellian content of a perceptual experience is a structured proposition involving the attribution of certain properties to an object. Such a structured proposition determines a secondary intension, but it is not itself a secondary intension. Likewise, one might hold that the Fregean content of a perceptual experience is a structured complex made up of the Fregean contents of its associated aspects: the Fregean content associated with the color, and perhaps those associated with the object, the location, and other aspects of the experience. We might see this content as a structured intension, made up of the intensions associated with each aspect. This intension determines a function from centered worlds to truth-values, but it is not itself such a function.

It may also be that some specific aspects of experience have contents that are more fine-grained than intensions. For example, following Peacocke (1992), one might argue that when one sees an object as a square and then as a regular diamond, there is a phenomenological difference that corresponds to a difference in content (representing something as a square or a regular diamond) that cannot easily be modeled as a function from worlds to extensions. And perhaps one could make the case that there are some experiences that could not be veridical in any possible world: if so, their associated intension might be false in all worlds, but they might still intuitively have some nontrivial Fregean and Russellian contents. In these cases one might appeal to conditions on extension that are more fine-grained than functions from worlds to extensions (allowing substantial conditions that are impossible to satisfy, for example). Still, as before, these contents will at least determine associated intensions. In what follows the differences will not play a large role, so it will be useful to use primary and secondary intensions to analyze Fregean and Russellian contents.

In the two-dimensional analysis of belief and language, as I understand it, the two dimensions of content correspond to the evaluation of epistemic possibilities (ways the world might actually be, for all we know a priori) and subjunctive possibilities (ways the world might have been, but is not). The same goes for the analysis of perceptual experience.

Intuitively, a perceptual experience places a constraint on epistemically possible states of the world. For all I know a priori, there are many ways that the world could turn out to be. We can think of these ways as epistemic possibilities (in a broad sense), and we can model them using centered worlds. The same goes for the epistemic possibilities that confront a perceiver. For example, when I have a phenomenally red experience, it is epistemically possible that the object I am looking at has property P1 and that P1 normally causes phenomenally red experiences, and it is epistemically possible that the object I am looking at has property P2 and that P2 normally causes phenomenally red experiences. These two hypotheses correspond to different centered worlds W_1 and W_2 . Intuitively, whether W_1 or W_2 turns out to be actual, my experience will be veridical. So we can say that W_1 and W_2 both verify the experience. On the other hand, there are worlds W_3 and W_4 where the object that the subject at the center is looking at has P1 and P2 (respectively), but in which that property normally causes phenomenally green experiences. Intuitively, if W_3 or W_4 turn out to be actual, my experience will be falsidical. So we can say that W_3 and W_4 both falsify the experience. These intuitions can be encapsulated in the claim that the primary intension of the experience is true in W_1 and W_2 , but false in W_3 and W_4 .

We can likewise evaluate perceptual experiences with respect to counterfactual circumstances, considered via subjunctive conditionals. If I have a phenomenally red experience directed at a red book B, then I can ask: if my eyes had been closed, but the book had still been present, would things have been as I (actually) perceive them to be? The intuitive answer is yes. Or I can ask: if the book had been present with the same intrinsic properties, but if I had been such that the book normally caused phenomenally green experiences in me, would things have been as I (actually) perceive them to be? The intuitions about this case are a bit less clear, but there is at least some intuition that the answer is yes. Intuitions of this sort suggest that with respect to counterfactual circumstances, the perceptual experience will be veridical (at least with respect to color) roughly iff the book B is (physically) red in those circumstances. This can be encapsulated by saying that the secondary intension of the experience is true in a world iff book B is red in that world.

One might even define the secondary intension of a perceptual experience in terms of these subjunctive conditionals. One can say that the secondary intension of an experience E is

true at a world W iff: had W obtained, things would have been the way that they appear to be to the subject (actually) undergoing E . To the extent that our judgments about these subjunctive conditionals are not fully determinate, the corresponding secondary intension may not be fully determinate, but in any case the intension will capture some of our intuitions about content.

Defining the primary intension of a perceptual experience is not as straightforward. So far I have talked in an intuitive way of a condition on extension, and about the evaluation of epistemic possibilities, but this falls short of a definition. One could simply leave a notion here as basic, but it would be nice to do more. One thought is to appeal to indicative conditionals. For example, one could hold that the primary intension of an experience E is true in a centered world W iff: if the actual world is qualitatively like W (and if I am in the position of the person at the center of W at the time of W), then E is veridical. The trouble is that one may want to evaluate the experience at worlds where the experience is not itself present. It is not obviously part of the content of an experience E that E itself obtains: one may want to hold that the content largely constrains the external world. Furthermore, in the case of beliefs and utterances, there are good reasons not to define primary intensions in terms of conditionals that make explicit reference to whether a belief or utterance would be true as it occurs in some circumstance (see Chalmers 2004b), and these reasons plausibly extend to the case of perceptual experience. So it would be useful to have a definition that applies more broadly than this.

In the case of beliefs, one can define a primary intension in terms of the belief's inferential connections. In particular, one can say that the primary intension of belief B is true at a centered world W if there is an appropriate inferential connection between the hypothesis that W is actual and B . (The hypothesis that W is actual can be understood more strictly as the hypothesis that D is the case, where D is an appropriate canonical description of W .) One way of understanding the inferential connection is in terms of rational inference: if a subject were to accept that W is actual, ought they rationally to accept B ? Another, arguably better way, is to understand it in terms of a priori entailment: if the subject accepts that W is actual, ought (idealized) a priori reasoning from there lead them to accept B ?

A definition along these lines is not as straightforward in the case of perceptual experiences, as the notion of an inferential connection from an arbitrary hypothesis to a perceptual experience is not entirely clear. However, one can give such a definition by relying on certain inferential connections between perceptual experiences and beliefs that are somewhat clearer. In particular, we have a reasonable grip on what it is for a subject to *take an experience at face value*, yielding a perceptual belief. In such a case, we can say that the

perceptual belief *endorses* the perceptual experience. Note that a belief that endorses a perceptual experience should be distinguished from a belief that the perceptual experience is true. The latter is a belief directed at the experience, but the former is a belief directed at the world.

The notion of endorsement is a primitive in the current account, but one can say some things to characterize it. Endorsement is a sort of truth-preserving inference between perception and belief: when a belief endorses a veridical experience, the belief will be true. This need not be true in reverse: that is, it can happen that a belief endorses a falsidical experience without itself being false. This can happen for a complex experience, for example, when intuitively the belief takes certain aspects of the experience at face value, and abstracts away from others. For example, a belief might endorse my current visual experience where color is concerned, and abstract away from shape. If my experience is veridical with respect to color but not to shape, the experience will be overall falsidical, but the endorsing belief will be veridical. On this conception, an endorsing belief may make fewer commitments than the experience that it endorses, but it cannot make any commitments that are not made by the experience that it endorses.

One can also introduce a notion of *complete endorsement*: a belief completely endorses an experience when it endorses that experience in all its aspects. Complete endorsement preserves both truth and falsity: when a belief endorses a falsidical experience, the belief will be false. For experiences of any complexity, it is not clear that we have the capacity to completely endorse them (as it is not clear that we can have beliefs with the relevant detail and complexity), but it seems that the notion is reasonable at least for simple experiences, and as an idealization in the case of complex experiences.

With these notions in place, we can use them in conjunction with the already defined notion of the primary intension of a belief to characterize the primary intension of a perceptual experience. One can say that the primary intension of an experience E is true at a centered world W iff the primary intension of any possible belief that endorses E is true at W. One could also appeal to complete endorsement here, simply equating the primary intension of a perceptual experience with that of a belief that completely endorses it, but the definition in terms of endorsement has fewer commitments. It does not require that complete endorsements exist for all perceptual experiences. Intuitively, it simply requires that for any aspect of a perceptual experience, there is some belief that endorses that aspect.

On the two-stage view, which has two different standards of veridicality for perceptual experience, this picture will be complicated slightly. We could say that just as there are two standards of veridicality for perceptual experiences, there are two standards of truth for

perceptual beliefs: perfect and imperfect truth, say. On this model, endorsing a perfectly veridical experience will produce a perfectly true perceptual belief, and endorsing an imperfectly veridical experience will produce an imperfectly true belief. One could then associate a perceptual belief with two different primary intensions, corresponding to the standards of perfect and imperfect truth. These intensions will then ground two different primary intensions for an experience, corresponding to perfect and imperfect veridicality.

Alternatively, and perhaps more plausibly, one could say that there is just one standard of truth for beliefs, and two sorts of endorsement: perfect and imperfect endorsement, say. Intuitively, perfectly endorsing an experience requires holding that experience to the standard of perfect veridicality, while imperfectly endorsing it requires holding it only to the standard of imperfect veridicality. On this model, imperfectly endorsing an imperfectly (or perfectly) veridical experience will produce a true belief, perfectly endorsing a perfectly veridical experience will produce a true belief, but perfectly endorsing an imperfectly veridical experience may produce a false belief. On this model a belief will be associated with just one primary intension, but the two different endorsement relations will yield two different primary intensions for an experience. I will not choose between the two models here.

One might argue that there are representational aspects of perceptual experience that cannot be endorsed by any possible belief. For example, one might hold that extremely fleeting experiences, or experiences that are far outside attention, cannot be endorsed. Or one might argue that experiences in animals that lack concepts cannot be endorsed. I think in these cases one could make a case that it is at least *possible* for the experiences to be endorsed, perhaps assuming some idealization of the subject's actual cognitive capacity. But this is clearly a substantive issue and the answer is unclear. For this reason, the characterization above is probably best not regarded as a definition, but it is at least a useful characterization that gives reasonably precise results for a wide range of ordinary experiences.

A subtlety arises concerning the object-oriented aspect of the experience. If one holds (not implausibly) that the Russellian content of a non-hallucinatory experience involves the object of the experience, then there should be a corresponding element of the Fregean content of the experience that picks the object out. Here the natural suggestion is this: Fregean content picks out that object as the object that is causing (or that is appropriately perceptually connected to) the current experience. This suggests a primary intension that maps a centered world to the object causing (or perceptually connected to) the relevant experience of the subject at the center of the centered world. To model this, one may need to build in a marked experience to the center of the relevant centered worlds, as the subject at the center may have many experiences (this is something that one has to do in any case; see Chalmers 2002a,

section 9). One can then say that the primary intension of a phenomenally red experience is true in a centered world if the object that causes (or is appropriately connected to) the marked experience at the center of that world has the property that normally causes phenomenally red experiences

Once again, this characterization of the relevant intension is just a first approximation. A detailed account would require consideration of the complexities of color experience, and also would require an analysis of the Fregean contents of the experience of spatial location and perhaps of the object-oriented aspect of the experience. For example, one might hold that where an experience has a Russellian content involving a particular object, there is an associated Fregean content that picks out that object as the object that is causing the current experience. This would require a primary intension that maps a centered world to the object causing (or otherwise perceptually connected to) a relevant experience of the subject at the center of the centered world (see section 10.2 here).

I should note that there is no requirement that either Russellian content or Fregean content be *conceptual* content. That is, it is not required that to have an experience with a given Russellian or Fregean content, a subject must have concepts with corresponding content. Frege's own favored variety of content was a sort of conceptual content: Fregean senses were always grasped through the possession of corresponding concepts. But the Fregean contents discussed here need not be Fregean in that respect. Rather, they are Fregean in the sense that they involve modes of presentation of objects and properties in the world, whether or not these modes of presentation involve concepts.

My own view is that a subject can have perceptual experiences with quite determinate content even without possessing corresponding concepts. For example, one can have a color experience that represents a given shade even without possessing a concept of that shade. It is nevertheless plausible that such an experience possesses a Russellian content (attributing that shade to an object), and it also possesses a Fregean content (one that is true if the object appropriately related to the perceiver has the property that normally causes the relevant sort of experience). If so, these contents are nonconceptual contents: an experience can have such contents even if the subject lacks a concept of the relevant shade, lacks a concept of the relevant experience, and lacks a concept of "normally causes". I will not argue for this claim here, however, and it is inessential for the other claims in this chapter.

These nonconceptual Fregean contents need not themselves be a different sort of object from conceptual Fregean contents. For example, it could be that an experience as of a red object and a belief that completely endorses it have the same Fregean content. Instead, one might say that the Fregean content of perception involves a nonconceptual content *relation*:

the relation that associates perceptual states with their Fregean contents is such that subjects need not possess the relevant concepts in order for their states to have the relevant content.²³ A pluralist about content can hold that there are both conceptual and nonconceptual content relations, but nonconceptual content relations are likely to be particularly useful in the analysis of perception.

Bibliography

- Block, Ned. 1990. Inverted Earth, *Philosophical Perspectives*. 4:53-79.
- Boghossian, Paul and Velleman, J. David. 1989. Color as a Secondary Quality. *Mind*. 98:81-103.
- Burge, Tyler. 1991. Vision and Intentional Content. In Ernie Lepore and Robert van Gulick, eds. *John Searle and his Critics*. Oxford: Blackwell.
- Byrne, Alex. 2001. Intentionalism Defended, *Philosophical Review*, 110:199-240.
- Byrne, Alex and Hilbert, David. 2003. Color Realism and Color Science. *Behavioral and Brain Sciences*. 26:3-21.
- Byrne, Alex. forthcoming. Color and the Mind-Body Problem, *Dialectica*.
- Campbell, John. 1993. A Simple View of Color, in John Haldane and Crispin Wright, eds. *Reality, Representation, and Projection*. Oxford: Oxford University Press.
- Campbell, John. 2002. *Reference and Consciousness*. Oxford: Oxford University Press.
- Chalmers, David J. 1996. *The Conscious Mind: In Search of a Fundamental Theory*. Oxford: Oxford University Press.
- Chalmers, David J. 2002a. The Components of Content, in *Philosophy of Mind: Classical and Contemporary Readings*. New York: Oxford University Press.
- Chalmers, David J. 2002b. Does Conceivability Entail Possibility? In Tamar Gendler and John Hawthorne, eds. *Conceivability and Possibility*. Oxford: Oxford University Press.
- Chalmers, David J. 2003a. The Content and Epistemology of Phenomenal Belief. In Quentin Smith and Alexandr Jokic, eds. *Consciousness: New Philosophical Perspectives*. Oxford: Oxford University Press.
- Chalmers, David J. 2003b. The Matrix as Metaphysics. Philosophy section of *The Matrix* website. [<http://consc.net/papers/matrix.html>]
- Chalmers, David J. 2004a. The Representational Character of Experience. In Brian Leiter, ed. *The Future for Philosophy*. Oxford: Oxford University Press.
- Chalmers, David J. 2004b. Epistemic Two-Dimensional Semantics. *Philosophical Studies*. 118:153-226.

²³ Heck (2000) distinguishes the “content view” of nonconceptual content (on which nonconceptual contents are objects quite distinct from conceptual contents) from the “state view” (on which nonconceptual contents are contents of nonconceptual states). One might call the conception suggested here the “relational view” of nonconceptual content. It is somewhat closer to the state view than to the content view, but it need not invoke the idea of a “conceptual state”.

- Chisholm, Roderick. 1942. The Problem of the Speckled Hen. *Mind*. 204:368-73.
- Clark, Austen. 2000. *A Theory of Sentience*. Oxford: Oxford University Press.
- Dretske, Fred. 1995. *Naturalizing the Mind*. Cambridge, Mass.: MIT Press.
- Edwards, James. 1998. The Simple Theory of Colour and the Transparency of Sensory Experience. In Cynthia MacDonald, Barry Smith, and Crispin Wright, eds. *Knowing Our Own Minds: Essays on Self-Knowledge*. Oxford: Oxford University Press.
- Egan, Andy. forthcoming. Appearance Properties?, *Nous*.
- Hardin, C.L. 1987. *Color for Philosophers*. Indianapolis: Hackett Publishing Company.
- Harman, Gilbert. 1990. The Intrinsic Quality of Experience, *Philosophical Perspectives*. 4:31-52.
- Heck, Richard. 2000. Nonconceptual Content and the Space of Reasons. *Philosophical Review*. 109:483-523.
- Holman, Emmett. 2002. Color Eliminativism and Color Experience. *Pacific Philosophical Quarterly*. 83:38-56.
- Horgan, Terence and Tienson, John. 2002. The Intentionality of Phenomenology and the Phenomenology of Intentionality. In David J. Chalmers, ed. *Philosophy of Mind: Classical and Contemporary Readings*. New York: Oxford University Press.
- Jackson, Frank. 1996. The Primary Quality of Color. *Philosophical Perspectives*. 10:199-219.
- Johnston, Mark. 1992. How to Speak of the Colors. *Philosophical Studies*. 68:221-63.
- Johnston, Mark. 2004. The Obscure Object of Hallucination. *Philosophical Studies*. 120:113-83.
- Johnston, Mark. forthcoming. *The Manifest*.
- Kriegel, Uriah. 2002. Phenomenal Content, *Erkenntnis*. 57:175-98.
- Levine, Joseph. 2003. Experience and Representation, in Quentin Smith and Aleksander Jokic, eds. *Consciousness: New Philosophical Essays*. Oxford: Oxford University Press.
- Lycan, William G.. 1996. *Consciousness and Experience*. Cambridge, Mass.: MIT Press.
- Lycan, William G. 2001. The Case for Phenomenal Externalism, *Philosophical Perspectives*. 15:17-35.
- Martin, M.G.F. 2002. Particular Thoughts and Singular Thought. In Anthony O'Hear, ed. *Logic, Thought, and Language*. Cambridge: Cambridge University Press.
- Maud, J.B. 1995. *Colours: Their Nature and Representation*. Cambridge: Cambridge University Press.
- McGinn, Colin. 1996. Another Look at Color, *Journal of Philosophy*. 93:537-53.
- McLaughlin, Brian. 2003. Color, Consciousness, and Color Consciousness. In Quentin Smith and Aleksander Jokic, eds. *Consciousness: New Philosophical Essays*. Oxford: Oxford University Press.
- Peacocke, Christopher. 1983. *Sense and Content: Experience, Thought, and Their Relations*. Oxford: Oxford University Press.
- Peacocke, Christopher. 1992. *A Study of Concepts*. Oxford: Oxford University Press.
- Putnam, Hilary. 1975. The Meaning of 'Meaning'. In *Mind, Language, and Reality*. Cambridge: Cambridge University Press.

- Searle, John R. 1983. *Intentionality*. Cambridge: Cambridge University Press.
- Searle, John R. 1990. Consciousness, Explanatory Inversion, and Cognitive Science. *Behavioral and Brain Sciences*. 13:585-595.
- Shoemaker, Sydney. 1990. Qualities and Qualia: What's in the Mind? *Philosophy and Phenomenological Research*. supplement, 50:109-131.
- Shoemaker, Sydney. 1994. Phenomenal Character. *Nous*. 28:21-38.
- Shoemaker, Sydney. 2001. Introspection and Phenomenal Character. *Philosophical Topics*. Reprinted in David J. Chalmers, ed. *Philosophy of Mind: Classical and Contemporary Readings*. Oxford: Oxford University Press.
- Shoemaker, Sydney. forthcoming. On the Way Things Appear. In Tamar Gendler and John Hawthorne, eds. *Perceptual Experience*. Oxford University Press.
- Siegel, Susanna. forthcoming. 'Which Properties are Represented in Perception?' In Tamar Gendler and John Hawthorne, eds. *Perceptual Experience*. Oxford University Press.
- Siegel, Susanna. forthcoming. Subject and Object in the Contents of Visual Experience.
- Siewert, Charles. 1998. *The Significance of Consciousness*. Princeton: Princeton University Press.
- Stoljar, Daniel. forthcoming. Consequences of Intentionalism. *Philosophical Studies*.
- Thau, Michael. 2002. *Consciousness and Cognition*. New York: Oxford University Press.
- Thompson, Brad. 2003. *The Nature of Phenomenal Content*. Ph.D. dissertation, University of Arizona.
- Thompson, Brad. forthcoming. The Spatial Content of Visual Experience. faculty.smu.edu/bthomps/spatialcontent.html.
- Tye, Michael. 1995. *Ten Problems of Consciousness: A Representational Theory of the Phenomenal Mind*. Cambridge, Mass.: MIT Press.
- Tye, Michael. 2003. Blurry Images, Double Vision, and Other Oddities: New Problems for Representationalism? In Quentin Smith and Aleksander Jokic, eds. *Consciousness: New Philosophical Essays*. Oxford: Oxford University Press.
- Wright, Wayne. 2003. Projectivist Representationalism and Color. *Philosophical Psychology*. 16:515-33.