

The Representational Character of Experience

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

1 Introduction

Consciousness and intentionality are perhaps the two central phenomena in the philosophy of mind. Human beings are conscious beings: there is something it is like to be us. Human beings are intentional beings: we represent what is going on in the world. Correspondingly, our specific mental states, such as perceptions and thoughts, very often have a phenomenal character: there is something it is like to be in them. And these mental states very often have intentional content: they serve to represent the world.

On the face of it, consciousness and intentionality are intimately connected. Our most important conscious mental states are intentional states: conscious experiences often inform us about the state of the world. And our most important intentional mental states are conscious states: there is often something it is like to represent the external world. It is natural to think that a satisfactory account of consciousness must respect its intentional structure, and that a satisfactory account of intentionality must respect its phenomenological character.

With this in mind, it is surprising that in the last few decades, the philosophical study of consciousness and intentionality has often proceeded in two independent streams. This was not always the case. In the work of philosophers from Descartes and Locke to Brentano and Husserl, consciousness and intentionality were typically analyzed in a single package. But in the second half of the twentieth century, the dominant tendency was to concentrate on one

Published in (B. Leiter, ed) *The Future for Philosophy*. Oxford University Press, 2004. I have presented versions of this paper, starting in 2001, at MIT, Cornell, Rutgers, Brown, and at the 2002 NEH Summer Institute on Consciousness and Intentionality. I would like to thank audiences on all of these occasions for discussion. Thanks especially to all the participants in the NEH Institute for many valuable conversations, and to Torin Alter, Justin Fisher, Terry Horgan, Uriah Kriegel, Brian Leiter, David Pitt, Sydney Shoemaker, Susanna Siegel, Brad Thompson, Leora Weitzman, and Wayne Wright for comments on a draft of this paper. Brad Thompson's dissertation on phenomenal content explores many of the same issues as this paper, and I have had much valuable discussion with him. Finally, this work has been influenced by the work of others in various obvious ways, but it has been influenced by the work of Sydney Shoemaker and Charles Siewert in ways that may be deeper than the obvious.

topic or the other, and to offer quite separate analyses of the two. On this approach, the connections between consciousness and intentionality receded into the background.

In the last few years, this has begun to change. The interface between consciousness and intentionality has received increasing attention on a number of fronts. This attention has focused on such topics as the representational content of perceptual experience, the higher-order representation of conscious states, and the phenomenology of thinking. Two distinct philosophical groups have begun to emerge. One group focuses on ways in which consciousness might be grounded in intentionality. The other group focuses on ways in which intentionality might be grounded in consciousness.

Grounding consciousness in intentionality. Those who take this approach typically try to analyze consciousness in terms of intentionality, without remainder. Some (such as David Rosenthal and Peter Carruthers) analyze conscious states in terms of higher-order states that represent them. But perhaps the most popular approach (taken by Michael Tye and Fred Dretske, among others) focuses on the first-order intentional content of conscious states, and advocates *representationalism*, analyzing conscious states as a certain sort of first-order intentional state. This approach is often combined with the view that intentionality can be explained in physical terms, to motivate a physically acceptable explanation of consciousness.¹

Grounding intentionality in consciousness. Those who take this approach argue that phenomenology plays a crucial role in grounding representational content. Some (such as John Searle) argue that consciousness is the ground of all true intentional content. Others (such as Terence Horgan and John Tienson, Colin McGinn, and Charles Siewert) argue that there is at least a distinctive and crucial sort of intentional content that accrues in virtue of the phenomenal character of mental states. These theorists do not typically offer reductive analysis of intentionality in terms of consciousness, but they typically hold that consciousness has a certain priority in the constitution of intentionality.²

My own sympathies lie more with the second group, but I think that the ideas of both groups deserve close attention. In particular, I think that there is significant promise in representationalism. Representationalism is often understood as taking a reductive view of consciousness, but this association is inessential. I think that the insights of

¹ See Rosenthal 1997, Carruthers 2000, Tye 1995, and Dretske 1995. Others who take this sort of approach include Harman 1990 and Lycan 1996.

² See Searle 1990, Horgan and Tienson 2002, McGinn 1988, and Siewert 1998. Others who take this sort of approach include Strawson 1994 and Loar 2003.

representationalism cohere equally well with a view on which consciousness cannot be reduced to anything more basic than itself.

In what follows I proceed largely from first principles in analyzing aspects of the relationship between consciousness and intentionality, while leaning on the analyses of others along the way. I focus on the phenomenal character and the intentional content of perceptual states, canvassing various possible relations among them. I argue that there is a good case for a sort of representationalism, although this may not take the form that its advocates often suggest. By mapping out some landscape, I try to open up territory for different and promising forms of representationalism to be explored in the future. In particular, I argue for a nonreductive, narrow, and Fregean variety of representationalism, which contrasts strongly with more widely explored varieties. I conclude with some words about the fundamental relationship between consciousness and intentionality.

2 Phenomenal properties and representational properties

Consciousness involves the instantiation of *phenomenal properties*. These properties characterize aspects of what it is like to be a subject (what it is like to be me right now, for example, or what it is like to be a bat), or what it is like to be in a mental state (what it is like to see a certain shade of green, for example, or what it is like to feel a certain sharp pain). Whenever there is something it is like to be a subject, that subject has specific phenomenal properties. Whenever there is something it is like to be in a mental state, that state has specific phenomenal properties. For many purposes it does not make much difference whether one focuses on the phenomenal properties of subjects or of mental states (it is easy to translate between the two ways of talking), and in what follows I will move back and forth between them.

Intentionality involves the instantiation of *representational properties*. We can say that a *pure representational property* is the property of representing a certain intentional content (or the property of having a certain intentional content; I will usually not pay much attention to this grammatical distinction). Intuitively, this involves representing things as being a certain way in the world. A perceptual state might have the content that there is a green object in front of me, and a belief might have the content that javelinas live in the desert. We can be neutral on just what sort of content this is: it might involve a set of possible worlds, a complex of objects and properties, a complex of Fregean senses, or something else again. What is most important is that intentional contents have *conditions of satisfaction*: they are the sort of thing that can be satisfied or can fail to be satisfied by states of the world. As with phenomenal

properties, we can regard representational properties as being instantiated by either subjects or by mental states.

An *impure representational property* is the property of representing a certain intentional content in a certain manner. This involves representing such-and-such as being the case in such-and-such a way. Here the “way” is a *manner of representation*, and involves a mental characterization of the state of representing. There are many different manners of representation. For example, one can represent a content perceptually, and one can represent a content doxastically (in belief): these correspond to different manners of representation. At a more fine-grained level, one can represent a content visually or auditorily. Manners of representation may also involve functional characterizations of the representing state. For example, one can represent a certain content in such a way that the content either is or is not available for verbal report.³

Whenever an impure representational property is instantiated, the corresponding pure representational property will also be instantiated (whenever a content is represented in a certain way, it is represented). And whenever a pure representational property is instantiated, numerous corresponding impure representational properties will be instantiated (whenever a content is represented, it is represented in certain ways).

I will take *representationalism* to be the thesis that phenomenal properties are identical to certain representational properties (or that they are equivalent to certain representational properties; see below). We can say that *pure representationalism* is the thesis that phenomenal properties are identical to pure representational properties, while *impure representationalism* is the thesis that phenomenal properties are identical to impure representational properties. In practice, almost all representationalists are impure representationalists, for reasons to be given shortly.

Some representationalists, such as Dretske and Tye, occasionally put their view by saying that phenomenal properties are identical to certain *represented* external properties, such as physical redness. As I am putting things, that would be a category mistake: phenomenal properties are by definition properties of subjects or of mental states, and physical redness is not (or need not be). I think that this is simply a terminological difference, however. For example, Dretske defines phenomenal properties (“qualia”) as the properties we are directly aware of in perception, and concludes these are properties such as colors. This is quite compatible with the claim that phenomenal properties in my sense are representational

³ Manners of representation are closely related to what Crane (2003) calls “intentional modes”. I do not use this way of speaking as it invites confusion with the quite distinct notion of modes of presentation (which are contents rather than psychological properties), used later in this article.

properties, as long as one holds that one is directly aware of the represented property rather than the representational property. Once we make the relevant translation, I think that these representationalists' most important claims can be put in the terms used here without loss.

3 Relationships

What are the relationships between phenomenal and representational properties? Do phenomenal properties *entail* (pure or impure) representational properties, or vice versa? Are phenomenal properties *identical* to representational properties? Are phenomenal properties *reducible* to representational properties, or vice versa?

We can say that one property entails another property when it is necessarily the case that whatever has the first property has the second property. We can say that two properties are equivalent when each entails the other: that is, if it is necessarily the case that whatever has one of the properties has the other. I will take it in what follows that the equivalence of two properties suffices for them to be identical. Not much rests on this: if one thinks properties are finer-grained than this, one can simply replace talk of identity throughout this paper by talk of equivalence.

This approach yields the following more detailed characterization of representationalism. Representationalism (about a class of phenomenal properties) is the thesis that for every phenomenal property (in that class), there is some representational property such that necessarily, a mental state (or a subject) has that representational property if and only if it has that phenomenal property.

I will concentrate on the class of perceptual phenomenal properties, and in some places more specifically on the class of visual phenomenal properties. So in the first instance, I am interested in the relationship between visual phenomenal properties and (pure or impure) representational properties, though it is natural to hope that this relationship will generalize to other phenomenal properties. In what follows, I canvas various possible relations among these properties, one step at a time.

Do pure representational properties entail phenomenal properties? That is, is it necessarily the case that any subject (mental state) that represents a certain content also has a certain phenomenal character? Here, it is most natural to answer no. The reason is that it seems that most or all representational contents can be represented *unconsciously*, without any associated phenomenal character at all. There can be unconscious beliefs, and subconscious perception, and it seems that in principle these unconscious states can represent a very wide range of intentional contents. One could preserve the entailment thesis by

denying that there are unconscious representational states, but that view is now almost universally rejected.

There are some weaker theses that may be left on the table here. First, one could hold that there are some special intentional contents (perhaps concerning phenomenal properties themselves, or intimately related to phenomenal properties) that cannot be represented unconsciously. If so, the corresponding pure representational properties would entail phenomenal properties. Second, one could hold (with Searle 1990) that representing an intentional content requires that one have some corresponding conscious state at some time or other, or at least potentially. If so, there could be an entailment from pure representational properties to certain diachronic or dispositional phenomenal properties (the property of having certain phenomenal properties in the past or the future, or potentially), if not to phenomenal properties in the strict sense. Finally, one could hold that there is an entailment from certain *impure* representational properties to phenomenal properties. I will return to this possibility below.

Do phenomenal properties entail pure representational properties? That is, is it necessarily the case that any subject with a certain phenomenal character thereby represents the world as being a certain way? Here the answer is quite plausibly yes, at least in the case of perceptual phenomenal properties. It seems intuitively clear that perceptual phenomenology, by its very nature, involves the representation of the external world. For example, my visual phenomenology right now involves the representation of a computer on a desk in front of me, with various books and papers scattered on the desk.

The claim that phenomenal experiences (in a given class) always have representational content is sometimes called *weak representationalism*. Weak representationalism (at least about visual experiences) is extremely plausible, and rarely denied.⁴ A stronger thesis also seems plausible: that any two phenomenally identical states will share some aspect of their representational content. As Colin McGinn has put it, there is a strong intuition that something about representational content is *internal* to phenomenology. That is, given a specific phenomenology, it seems that if a mental state has this phenomenology, it must also have a certain specific representational content. If this is right, then phenomenal properties entail certain pure representational properties.

This line of thought has been developed into an argument at greatest length by Charles Siewert (1998). I will not recapitulate Siewert's extensive arguments here, but the basic idea

⁴ Weak representationalism is characterized in roughly this way by Crane 2003, Levine 2003 and Lycan 2000. Even the views of such anti-representationalists as Block and Peacocke appear to be compatible with weak representationalism.

is that visual experiences are *assessible for accuracy*, in virtue of their phenomenal character. For example, when I have a visual experience as of something X-shaped in front of me, this experience may be either accurate or inaccurate, depending on what is really in front of me. Further, it seems that any visual experience with the same phenomenal character would be assessible for accuracy in the same sort of way: roughly, it would be correct if and only if there is an object with an appropriate shape in front of the subject. If this is right, then a phenomenal property (having an experience with the phenomenal character as of seeing something X-shaped) entails a pure representational property (roughly, representing that there is an object with a certain shape in the world).

Something similar plausibly applies to other visual experiences, such as experiences of position and of color, and to other perceptual experiences, such as auditory and tactile experiences. Siewert argues that this reasoning can be extended to the phenomenology and intentionality of thought, as well as perception, but this raises complications that I will not go into here. For now, what matters is that there is plausibly an entailment from perceptual phenomenal properties to pure representational properties.

A similar point can be made for subjects of experience, as well as for mental states. Terence Horgan and John Tienson (2002) have argued that necessarily, any two subjects that are phenomenal duplicates will share significant representational content.⁵ Even a brain in a vat that has experiences that are phenomenally identical to mine will represent its environment as being in significant respects just as I represent my environment as being. This suggests that the phenomenal properties of a subject entail pure representational properties for that subject.

Just what sort of representational properties are entailed by phenomenal properties? Some tricky questions arise here. If phenomenal properties are intrinsic properties of an individual, then presumably the corresponding representational properties must be intrinsic too. One might reasonably wonder just how to characterize these properties, given the apparent extrinsicness of many representational properties. I will look more closely at these questions later. For now, I will take it that there is at least a strong *prima facie* case that phenomenal properties entail pure representational properties.

Are phenomenal properties identical to pure representational properties? That is, for a given phenomenal property, is there some representational content that is represented if and only if that phenomenal property is instantiated? The answer to this question is plausibly no, for reasons given above: identity requires entailment, and it is plausible that phenomenal

⁵ Loar (2003) makes a related argument from phenomenal duplicates, although his conception of phenomenal intentionality (unlike Horgan and Tienson's) is shorn of the connection to conditions of satisfaction.

properties are not entailed by pure representational properties. If any given representational content can be represented unconsciously, then pure representational properties cannot be identical to phenomenal properties, and pure representationalism is false.

The only hope for pure representationalism is to deny that any given representational content can be represented unconsciously. The pure representationalist might deny that unconscious representation is possible, or might hold that there are some special representational contents that cannot be represented unconsciously and can be represented only by a specific sort of phenomenal state.⁶ If such contents exist, the property of representing one of them might be identical to a phenomenal property. In the absence of strong reasons for thinking that such contents exist, however, I will pass over this option in what follows.

Are phenomenal properties identical to impure representational properties? We have seen that pure representational properties are plausibly entailed by phenomenal properties, but that the reverse is not the case, because of the possibility of unconscious representation. This rules out an identity between phenomenal properties and pure representational properties. However, it leaves open the possibility of an identity between phenomenal properties and *impure* representational properties, which require an appropriate manner of representation. If we stipulate the right manner of representation, the possibility of unconscious representation might be ruled out, removing at least one obstacle to an identity.

The obvious suggestion is to require a *phenomenal* manner of representation. Representational contents can be represented either phenomenally (as in conscious vision, for example) or nonphenomenally (as in unconscious thought). For a given content, we can define a corresponding impure representational property: the property of phenomenally representing that content. It is natural to wonder whether phenomenal properties might be identical to impure representational properties of this sort.

For a given phenomenal property, one can find a corresponding impure representational property of this sort. If the phenomenal property entails some representational properties (as at least perceptual phenomenal properties plausibly do), then it will entail some maximally specific representational property, involving a specific representational content. We can then define the impure representational property of phenomenally representing that content. The original phenomenal property will clearly entail this impure representational property. If the impure representational property itself entails the original phenomenal property, then the two properties are identical.

⁶ Something like pure representationalism is advocated by Lloyd 1991, who takes the first strategy, and by Thau 2002, who takes the second strategy.

The only obstacle to this identification is the possibility that two distinct phenomenal properties might correspond to the same impure representational property. This will happen if two distinct phenomenal states can have exactly the same representational content. This would happen if a perceptual state and a nonperceptual state (a visual experience and a conscious belief, for example) could phenomenally represent the same content. It would also happen if two perceptual experiences in different modalities (a visual experience and an auditory experience) could have the same content. Finally, it would happen if two phenomenally distinct perceptual experiences in the same modality (two visual experiences, for example) could have the same content.

It is not clear whether any of these cases can arise. But if they arise, there is a natural strategy for dealing with them: one can add more specificity to the manner of representation. Taking the first sort of case, one might argue that an experience of color and a belief about color could in principle have the same content, perhaps representing an object as having a certain specific shade of red. If this sort of case can arise, the perceptual phenomenal property will not be identical to the property of phenomenally representing the content in question. But one can handle the case by moving to a more specific impure representational property, such as the property of *perceptually* phenomenally representing the content in question.

Likewise, one might argue that a visual experience and an auditory experience might have the same content, perhaps representing that there is an object on one's left. In this case, it is far from clear that such states can have the same *specific* representational content, which is what is needed for a counterexample to equivalence. In most of the cases one might introduce, even if there is an equivalence in unspecific content (that there is an object on one's left), there is usually a difference in specific content (say, that there is a red spherical object on one's left, or that there is an object making a certain sound on one's left). So it is far from clear that this sort of case can arise. But if it can arise, one need only move once again to a more specific impure representational property, such as the property of *visually* (or auditorily) phenomenally representing the content in question.

Finally, one might argue that two phenomenally distinct visual experiences might have the same content. It is not easy to see how this could happen. Byrne (2001) has argued that this could not happen, on the grounds that any phenomenal difference between visual experiences is a difference in how things seem, and any difference between how things seem is a difference in representational content. Even if one resists this argument (on the grounds, perhaps, that it involves an equivocation on "how things seem"), it is not easy to come up with specific cases of phenomenally distinct visual experiences with the same content. When

we conceive of phenomenally distinct visual experiences, it seems that we almost always conceive of experiences that differ in aspects of their representational content.

Peacocke (1983) has argued that there are such cases: for example, images of two trees at different distances might represent them as having the same height, while the images take up differently-sized portions of the visual field. But representationalists (e.g. Tye 1995) have plausibly replied that these images differ in certain representational properties (such as their representation of distance, or of visual angle). Block (1990) has argued that subjects with phenomenally distinct color experiences in different environments might represent the same external colors. But one can argue, along lines developed later in this paper, that such experiences always differ in certain key aspects of their representational content. Other purported examples, involving cases of blurry vision, double images, and the like, have received reasonably plausible replies (e.g. Tye 2003) from representationalists.

To my mind, the most plausible potential cases of phenomenally distinct visual experiences with the same representational content involve differences in attention. Shifts in attention clearly make a phenomenal difference to visual experiences. In typical cases, they also make a representational difference: for example, shifting attention to a word may lead one to represent the shapes of its letters with greater specificity. But there are cases that are less clear. For example, one might look at two red pinpoint lights against a black background, and shift attention from one to the other. Here it is not obvious that there is a representational difference between the cases. There are various suggestions one might make in response, however. One might argue that the position or color of the light to which one is attending is represented with greater specificity than that of the light to which one is not attending. Or one might argue that the light to which one is attending is represented as being more salient than the other light. Here it is not completely clear what sort of property “salience” is, but it is plausible that there is such a property (though it may be relative to a subject at a time), and the suggestion that it is represented in attention seems reasonably apt to the phenomenology.

If there are any cases of phenomenally different visual experiences with the same representational content, one will have to build further differences into the manners of representation in impure representational states (so one might have the property of attentively representing a certain content, and so on). But it is not clear that such cases exist. My own view is that the claim that phenomenal differences between visual experiences always correspond to representational differences has some *prima facie* plausibility, and serves as a sort of null hypothesis that should be rejected only if there is strong evidence against it. If the view is correct, then visual phenomenal properties are identical to impure representational properties of the form involving the (visual) phenomenal representation of a certain content.

Is representationalism correct? If what has gone before is correct, then there is a strong prima facie case for the truth of impure representationalism, at least about visual phenomenal properties. These properties are plausibly equivalent to certain impure representational properties. These impure representational properties have the character: *phenomenally representing such-and-such content*, or (if nonperceptual phenomenal states can have the same content) *perceptually phenomenally representing such-and-such content*, or (if experiences in a different perceptual modality can have the same content) *visually phenomenally representing such-and-such content*.

4 Reductive and nonreductive representationalism

I have argued that phenomenal properties are plausibly identical to certain impure representational properties. In the characterization of these impure representational properties, however, the concept of the *phenomenal* itself plays an important role. So the analysis I have given certainly does not offer any sort of reductive analysis of phenomenal properties in terms of nonphenomenal notions.

Here, it is useful to distinguish *reductive* from *nonreductive* representationalism. Representationalism is the thesis that phenomenal properties are identical to certain pure or impure representational properties. Reductive representationalism holds that phenomenal properties are identical to certain pure or impure representational properties that can be understood without appeal to phenomenal notions. Nonreductive representationalism holds that phenomenal properties are identical to certain pure or impure representational properties, where these cannot be understood without appeal to phenomenal notions.

In a little more detail: representationalism holds that any given phenomenal property is identical to the property of representing a certain content in a certain manner (or in the case of pure representationalism, to the property of representing a certain content). Here, the right-hand side invokes a number of notions: the notion of representation, the notion of a certain content, and the notion of a certain manner of representation. If all of these notions can be understood without appeal to phenomenal notions, then the right-hand side offers a reductive account of phenomenal properties. If any these notions can only be understood in terms of phenomenal notions, then the right-hand side offers a nonreductive account of phenomenal properties.⁷

⁷ One can also distinguish *epistemically* reductive representationalism from *metaphysically* reductive representationalism, depending on whether the identity between phenomenal properties and non-phenomenally-specified representational properties is held to be a priori (or conceptually true), or simply to be true.

The analysis I have given above is clearly not a reductive analysis. The manner of representation invoked in this analysis makes explicit appeal to the notion of the phenomenal. It may also be, for all I have said so far, that the contents involved in characterizing the representational properties must themselves be understood in terms of phenomenal notions (for example, if they involve the attribution of phenomenal properties to objects, or if they involve the attribution of certain dispositions to cause phenomenal properties). And it might even be that the notion of representation cannot be understood in wholly nonphenomenal terms. The status of the last two claims is unclear, but certainly the first claim makes it clear that this is not a reductive analysis.

Proponents of reductive representationalism need to do three things. First, they need to argue that an account of representation can be given in non-phenomenal terms. Second, they need to argue that the relevant contents can be specified in a manner that is independent of phenomenal notions (this is a version of what Peacocke (1983) calls the “Adequacy Constraint”). And third, they need to argue that the appeal to a phenomenal manner of representation can be replaced by some non-phenomenally characterized manner of representation.

In practice, this is just what reductive representationalists (such as Tye 1995 and Dretske 1995) try to do.⁸ They appeal to an independently motivated account of representation in causal, informational, and teleological terms. They argue that the relevant representational contents can be understood as involving the attributions of ordinary physical properties (such as surface reflectance properties, in the case of color experiences) to physical objects in the world, without any need for phenomenal notions. And they argue that the relevant manner of representation can be understood in wholly functional terms.

I will set aside the first two points for now, and focus on the third. For the reasons given earlier, it seems that a manner of representation is needed to distinguish phenomenal from nonphenomenal representation of a given content. At this point, Tye introduces the constraint that the representing state must be *poised* to play a certain role in the control of speech and action. This constraint is introduced precisely to distinguish phenomenal from nonphenomenal representation, in cases such as blindsight. Dretske’s account is less explicit about how to handle unconscious representation, but at some places he appeals to the

Epistemically reductive representationalism entails metaphysically reductive representationalism, but not vice versa. For example, Tye’s representationalism is metaphysically but not epistemically reductive.

⁸ Reductive representationalists include Clark 2000, Dretske 1995, Droege 2003, Harman 1990, Jackson 2003, Lycan 1996, and Tye 1995.

constraint that the representational states must play an appropriate role in the formation of beliefs.

These accounts involve broadly *functional* accounts of the manner of representation: they require that the relevant representational states play the right sort of functional role within the cognitive system. In effect, this functional constraint is offered as an account of the difference between phenomenal and nonphenomenal representations (of the content in question). To this extent, then, the accounts involve functionalism about the phenomenal/nonphenomenal distinction. We might say these accounts are a species of *functionalist representationalism*. They are not functionalist across the board (both Dretske and Tye lean on nonfunctionalist teleological accounts of what it is to represent the contents in question), but they involve functionalism about the crucial distinction between phenomenal and nonphenomenal representation of the relevant contents.

Functionalist representationalism might be seen as the result of *conjoining* the relatively neutral version of representationalism above with functionalism about the phenomenal/nonphenomenal distinction. If we grant that phenomenal properties are identical to the property of phenomenally representing a certain content, and if we grant that what it is to phenomenally represent a content is to represent the content with a state that plays an appropriate functional role, the result is just this sort of account. So if functionalism about the phenomenal/nonphenomenal distinction can be justified, then functionalist representationalism can be justified.

Conversely, however, if there is reason to doubt functionalism about the phenomenal/nonphenomenal distinction, then there is reason to doubt functionalist representationalism. And there are many reasons to doubt functionalism about the phenomenal/nonphenomenal distinction. While there is some reason to believe that phenomenality is *correlated* with certain functional properties (such as the property of global availability), there are familiar reasons to doubt that it is reducible to these properties. For example, it seems coherent to suppose that any such functional property can be instantiated without any associated phenomenology at all. Many have argued that such a situation is metaphysically possible, and some (e.g. Block 1979 and Searle 1981) have even argued that it is nomologically possible. Others (e.g. Levine 1983) have argued that there is an explanatory gap between a functional characterization of a state and a phenomenal characterization: there

is no clear reason why the fact that a state plays a certain functional role should somehow make it the case that there is something it is like to be in that state.⁹

I will not try to adjudicate these arguments here. I will simply note that the functionalism in these varieties of representationalism is entirely optional. It is also worth noting that this functionalism is largely unargued in the work of Dretske and Tye. Tye's defense consists mostly in noting that the functional criterion handles cases such as blindsight, but clearly this falls well short of establishing any sort of reduction; while Dretske is not explicit about the matter at all. In arguing for their views, both spend by far the greatest portion of their time arguing for representationalism, and the crucial functionalism is left mostly in the background.

This leaves the following situation: there are good arguments for representationalism, but these are not good arguments for reductive representationalism. There is good reason to believe that phenomenal properties are identical to impure representational properties involving the phenomenal representation of certain contents. Both the reductive and the nonreductive representationalist can accept this. The reductive representationalist must conjoin this with a reductive account of the phenomenal/nonphenomenal distinction. But there are familiar reasons to doubt that such an account is possible, and the arguments for representationalism give no special reason to think that such an account exists.¹⁰

For my part, I think that nonreductive representationalism is a much more plausible view.¹¹ This view articulates an internal connection between phenomenal properties and representational properties, without trying to reduce the phenomenal to a wholly nonphenomenal domain. At the very least, the nonreductive style of analysis above provides a reasonably neutral starting point for further representationalist analyses. Representationalists of many different stripes can agree that visual phenomenal properties are equivalent to impure

⁹ A teleological representationalist such as Dretske allows that any narrow functional property can be instantiated without phenomenology, as phenomenology requires the relevant sort of environment. However, the central anti-functional arguments (by Block, Searle, etc) also tend to suggest that any given functional property could in principle be instantiated in any given environment, without phenomenology. So "teleofunctionalist" views are vulnerable to the same sort of arguments as functionalist views.

¹⁰ Versions of this point are made by Chalmers 1996 (p. 377), Kriegel 2002a, Vinuesa 2000, and Warfield 1999. The point that representationalism on its own does not offer the prospect of closing the explanatory gap is also made in a slightly different way (focusing on the represented contents rather than the manner of representation) by Stoljar forthcoming.

¹¹ To my knowledge, nonreductive representationalism has not been discussed explicitly as such in the literature, but Byrne 2001 notes the possibility of representationalism without reduction, and Crane 2003 and Levine 2003 indicate sympathy for this combination. The views suggested by Horgan and Tienson 2002, Searle 1983, and Siewert 1998 seem broadly compatible with the same combination.

representational properties of the form: phenomenally (or visually phenomenally, or whatever) representing such-and-such a content.

The questions for further debate concern: (i) the specificity in the manner of representation (e.g. phenomenal, perceptual phenomenal, or visual phenomenal?); (ii) whether the phenomenal manner of representation can itself be reductively analyzed; and (iii) the nature of the represented contents. So far, I have addressed the first two questions. In what follows I will address the third.

5 Narrow and wide representationalism

It is widely believed that phenomenal properties depend only on the internal state of the subject. It is also widely believed (following the arguments of Putnam 1975 and Burge 1979) that most representational properties depend not only on the internal state of the subject, but also on the subject's environment. For example, it is widely accepted that if two physically identical subjects have relevantly different environments (e.g. an "Earth" environment containing H₂O, or a "Twin Earth" environment containing XYZ), the contents of their mental states may be relevantly different (e.g. one has beliefs about H₂O, while the other has beliefs about XYZ).

This poses a *prima facie* problem for representationalism. Let us say that a property is *narrow* when necessarily, for any individual who has the property, an intrinsic duplicate of that individual has the property (regardless of environment). Let us say that a property is *wide* when it is possible for an individual to have the property while an intrinsic duplicate lacks that property. Then the following three propositions form an inconsistent triad:

- (i) Phenomenal properties are equivalent to representational properties
- (ii) Phenomenal properties are narrow
- (iii) Representational properties are wide

One can react to this inconsistent triad in three different ways: one could deny (i), (ii), or (iii). The first strategy leads to *anti-representationalism*. Block (1990) takes this strategy, in effect using the plausibility of (ii) and (iii) as an argument against representationalism.

The second strategy leads to *wide representationalism*, on which both phenomenal properties and the representational properties that they are equivalent to are taken to depend on a subject's environment. This strategy is taken by Dretske (1995), Tye (1995), and Lycan (1996), who in effect use the plausibility of (i) and (iii) to argue against the narrowness of phenomenal properties.

Both of these strategies have significant costs. The anti-representationalist strategy has trouble doing justice to the persistent intuition that phenomenal duplicates are deeply similar in certain representational respects. The wide representationalist strategy is even more counterintuitive, entailing that what it is like to be a subject depends constitutively on factors that may be far away from the subject and in the distant past. Some wide representationalists (especially Lycan 2001) have tried to defend this view, but it has many odd consequences. To pick just one: a change in environment can often yield a gradual change from one wide representational content to another quite different content (e.g., from representing red to representing green), with a period of indeterminate representation of both contents (or divided reference) in the middle. But it is hard to know what an indeterminate phenomenal state involving both phenomenal red and phenomenal green could be like (striped? superimposed?), and it is hard to believe that a subject going through this change in environment would pass through such a state. Lycan suggests that there may be much less indeterminacy in perceptual representation than in cognitive representation, but it is hard to believe that indeterminacy can be eliminated entirely. If there is even a brief period of indeterminacy, then the puzzle arises.

I think that by far the best-motivated strategy is the third. This is the strategy of *narrow representationalism*, which holds that both phenomenal properties and the representational properties they are equivalent to depend only on a subject's internal state. The arguments by Putnam, Burge, and the like give good reason to think that *many* representational properties are wide, but they give little reason to think that *all* representational properties are wide. If we accept that some representational properties are narrow, the inconsistent triad disappears. The natural strategy is then to exploit narrow representational properties in developing a narrow representationalism.¹²

Narrow representational properties can most naturally be developed by an appeal to *narrow contents*: contents of the sort that, if they are represented by a subject, they are represented by any intrinsic duplicate of that subject (regardless of the environment). The arguments of Putnam and Burge suggest that many contents are wide: especially the contents of beliefs about natural kinds and about individuals, and those involving deference to a linguistic community. But these arguments are consistent with the view that many other contents are narrow.

For example, it is widely accepted that many logical and mathematical beliefs have narrow content. More generally, Horgan and Tienson (2002) have argued that the existence of

¹² Versions of narrow representationalism are advocated by Rey 1998, Kriegel 2002b, and Levine 2003.

rich narrow content is a consequence of the plausible view that phenomenal duplicates share rich representational content. And it is possible to hold on quite general grounds that all mental states have a sort of narrow content, in addition to any wide content they might have. On the latter sort of view, although two corresponding “water”-beliefs of subjects on Earth and Twin Earth may differ in that one is about H₂O and the other is about XYZ (wide content), the beliefs will also have significant narrow content in common, perhaps characterizing the relevant liquids in qualitative terms.

I have argued elsewhere (Chalmers 2002) for a view of the latter sort, holding that the narrow content of beliefs can be understood as a sort of Fregean epistemic content. I think that this sort of account can naturally be extended to give an account of the narrow content of perceptual states. But first, it is useful to examine different ways of characterizing representational content.

6 Russellian representationalism

One natural way of characterizing the content of mental states involves objects and properties in the world. A belief, such as the belief that Hesperus is bright, can be regarded as composed of concepts (which I will regard here as mental items of some sort). One can naturally say that concepts have *extensions*: objects and properties that are picked out by these concepts. The extension of the concept *Hesperus* is the planet Venus, and the extension of the concept *bright* is the property of brightness. One can then characterize the content of a belief in terms of the extensions of the concepts that are involved in the belief. For example, a belief that Hesperus is bright might have the content: [*Venus, brightness*]. Contents of this sort, involving complexes of specific objects and/or properties in the world, may be thought of as Russellian contents (following Russell, who thought all contents were like this).

Let us say that a Russellian representational property is the property of representing a certain Russellian content, or the property of representing a certain Russellian content in a certain way. Then *Russellian representationalism* holds that phenomenal properties are identical to Russellian representational properties. Most contemporary representationalists are Russellian representationalists.¹³

It is natural to think that perceptual experiences have Russellian content. They involve the attribution of properties to objects. For example, when I have a veridical experience of a green ball B, it is natural to think that I represent the ball as being green, and that my

¹³ Thompson 2003 also distinguishes Fregean and Russellian representationalism. McLaughlin 2003 uses “denotational intentionalism” for something like Russellian representationalism.

experience has an *object-involving* Russellian content roughly as follows: [*B, greenness*]. It is also natural to think that my experience has a sort of *existential* (or merely *property-involving*) Russellian content, to the effect that there is an object in front of me that is green.

To develop a Russellian representationalism, most Russellian representationalists invoke something like the property-involving Russellian content mentioned above, rather than an object-involving Russellian content. The reason is that it seems that phenomenology can be dissociated from the object-involving content. My visual experience when I look at a green ball has a certain object-involving content, but it seems that a phenomenally identical experience could lack that content, if I were looking at a different object, for example, or if I were hallucinating. (Some direct realists will deny this claim, but most contemporary representationalists accept it.) But in these cases, the property-involving content of the original experience is still plausibly present: my visual experience still represents that there is a green object in front of me.

To go into more detail, it is useful to focus on the case of phenomenal properties associated with colors (though I think the discussion that follows generalizes). Let us say that *phenomenal redness* is the phenomenal property typically associated in our community with seeing red things. (Phenomenal redness should be distinguished from ordinary redness: the first is a property of experiences or of subjects, the second is a property of objects in one's environment.) It is natural to think that instances of phenomenal redness have the content that certain objects in the environment are *red*: that is, these have contents that attribute the property of redness.

This line of thought naturally leads to the most common Russellian representationalist view. This view holds that phenomenal redness is equivalent to the impure representational property of (visually) phenomenally representing a certain Russellian content involving the attribution of redness. Of course, for this view to be made more specific, one needs to specify just what sort of property redness is.¹⁴

Most commonly, Russellian representationalists of this sort take the view that redness is a physical property, such as a certain surface spectral reflectance. Then phenomenal redness is equivalent to a certain sort of attribution of this physical property to objects. Something analogous is usually said about other perceptual phenomenal properties: all are equivalent to

¹⁴ The distinctions that follow between physical, projectivist, primitivist, and dispositionalist versions of representationalism are also made by Stoljar (forthcoming). Stoljar calls the first of these the "physicalist" interpretation of representationalism, but I avoid this usage to avoid a natural confusion. Stoljar suggests that projectivism is not really a form of representationalism because it appeals directly to qualia; but once we acknowledge the possibility of nonreductive representationalism, this problem falls away.

representational properties involving the attribution of physical properties. The resulting view might be called *physical Russellian representationalism*. This strategy has the advantage of being reductive, at least where the content is concerned: there is no need to use phenomenal notions in characterizing the relevant contents.

Physical Russellian representationalism is plausibly a sort of wide representationalism. To have the relevant representational property, the subject must represent the relevant physical property, which plausibly requires that the subject has been inhabiting an appropriate environment. This leads to numerous counterintuitive consequences. For example, it seems possible that a subject might have experiences phenomenally identical to my experiences of redness, even though their experiences have been caused by objects that are not physically red (i.e., objects that have a different surface spectral reflectance); and it seems possible that these color experiences in such a subject might nevertheless be veridical. But physical Russellian representationalism must deny this possibility. I will not pursue this matter here, as it has been explored at length by others (e.g. Block 1990; Shoemaker 2002).

If we are interested in the prospects for narrow representationalism, it seems that other forms of Russellian representationalism are more compatible with this view. For example, *projectivism* about colors holds that colors are phenomenal properties, or perhaps that they are qualitative properties of a visual field. *Projectivist Russellian representationalism* holds that color experiences attribute these properties to objects, and that phenomenal color properties are equivalent to corresponding representational properties.¹⁵ (One might embrace projectivist Russellian representationalism without embracing projectivism about colors, if one holds that the relevant contents do not involve attribution of colors.) This is most naturally seen as a sort of narrow and nonreductive representationalism: two intrinsic duplicates will share the relevant representational properties, given that they share their phenomenal properties.

A related view, *primitivism* about colors, holds that colors are certain primitive intrinsic properties that are not phenomenal properties or properties of a visual field, but that are nevertheless constitutively connected to such properties. On this view colors have an intrinsic “qualitative” nature that is revealed in some fashion by color experiences. *Primitivist Russellian representationalism* holds that color experiences attribute these primitive properties to objects, and that phenomenal color properties are equivalent to corresponding

¹⁵ Boghossian and Velleman 1989 have a view that seems compatible with projectivist Russellian representationalism: they advocate projectivism about colors, and hold that properties of this sort are represented in color experience.

representational properties.¹⁶ (One might embrace primitivist Russellian representationalism without embracing primitivism about colors, if one holds that the relevant contents do not involve attribution of colors.) This is again naturally seen as a sort of narrow representationalism, given that phenomenal duplicates represent the same primitive properties, and given that phenomenal properties are narrow. It is also plausibly seen as nonreductive, given that colors are taken to be irreducible properties that are constitutively tied to phenomenal properties.

There is a certain phenomenological plausibility in both projectivist and primitivist Russellian representationalism. Both involve the attribution of intrinsic properties to objects, properties with a qualitative nature that is closely connected to the qualities of experience itself. But these views seem to have the counterintuitive consequence that color experience is massively illusory. When we have an experience as of a red apple, it seems unlikely that the apple itself instantiates phenomenal redness, or that it instantiates a corresponding property of our visual field. It also seems unlikely that it instantiates a simple intrinsic property with a qualitative nature that is constitutively connected to the quality of my visual experience. Further: it seems plausible (as above) that subjects in different environments might have phenomenally red experiences that are caused by objects that share no relevant intrinsic properties, without either subject being in a privileged position with respect to veridicality. If both experiences attribute the same intrinsic properties, then both cannot be correct, so neither is correct. Generalizing, it seems that projectivist or primitivist representationalism lead to the consequence that all color experiences are illusory. Some have accepted this consequence, but it seems a high price to pay.¹⁷

Finally, *dispositionalism* about color holds that colors are dispositional properties, involving the disposition to cause certain sorts of experiences: for example, redness might be the disposition to cause phenomenally red experiences in a certain class of subjects in normal conditions. *Dispositionalist Russellian representationalism* holds that color experiences attribute these dispositional properties to objects, and that phenomenal color properties are equivalent to corresponding representational properties. As before, one might embrace dispositionalist Russellian representationalism without embracing dispositionalism about

¹⁶ The “figurative projectivism” of Shoemaker (1990) is a sort of primitivism, while projectivism corresponds to his “literal projectivism”. Advocates of (something like) primitivist Russellian representationalism include Maund 1995, Holman 2002, Jakab 2003, and Wright 2003 (who are also primitivists about colors) and Thau 2002 (who is not).

¹⁷ Boghossian and Velleman, Holman, Jakab, Maund, and Wright all accept that color experiences are illusory. Thau is not explicit about the matter.

colors, if one holds that the relevant contents attribute properties other than colors. This sort of view is most closely associated with Shoemaker (1994; 2001).

Dispositionalist representationalism characterizes the relevant representational contents in partly phenomenal terms, so it can easily be seen as a sort of nonreductive representationalism. Shoemaker does not see it this way. He conjoins the basic view with the view that phenomenal properties are identical to certain neurophysiological properties, and so holds that the relevant dispositional properties can be identified nonphenomenally. So his version of the view is a sort of reductive representationalism. But the reductive aspect of his view is largely independent of the representationalist aspect, and one might well embrace the basic shape of the representationalist view without embracing the reductive claim. If so, the result is a nonreductive version of representationalism.¹⁸

Dispositionalist representationalism can also easily be seen as a sort of narrow representationalism. It seems plausible to think that two intrinsic duplicates will share their representational contents: both will have phenomenally red experiences, attributing the disposition to cause red experiences to their objects. However (as Egan forthcoming and Thompson 2003 have pointed out), it is not clear just how this dispositional property should be characterized. If it is the disposition to cause phenomenally red experiences in a specific subject (the subject of the experience), the phenomenally red experiences in different subjects will attribute *different* dispositional properties to their objects, and will differ in the relevant content. This content will not be narrow (the attributed properties will differ between duplicate subjects); and worse, the corresponding representational property will not be entailed by the phenomenal property. An appeal to dispositions to cause the experience in a specific community has the same problem.

Partly to deal with this sort of problem, Shoemaker's version of dispositionalist representationalism appeals to the "higher-order" disposition to cause the experience in one or more sorts of creature in some circumstance; but this view has the consequence that illusory representation is impossible (every object will have the relevant disposition). Egan suggests that the experience attributes an *indexical* disposition: something like the disposition to cause phenomenally red experiences in *me*, where this disposition is attributed equally by both subjects; but it is not clear that this indexical disposition is a legitimate property, or that it could be shared by objects of veridical perception by different subjects in different environments. Alternatively, one might hold that one is attributing to *oneself* the property of being confronted by an object with the disposition to cause certain experiences in one; but it is

¹⁸ Kriegel 2002b and Levine 2003 embrace the basic shape of Shoemaker's dispositionalist representationalism, without Shoemaker's reductive claim.

counterintuitive to hold that experience involves the attribution of properties to oneself and not to an external object.

This is one difficulty with the view. Another difficulty is that it seems implausible for various reasons that colors are dispositions. It seems possible that in counterfactual circumstances, red objects might be disposed to cause phenomenally green experiences in normal observers, or in me. If this is right, redness is distinct from the disposition to cause phenomenally red experiences. It seems more compatible with our judgments about the cases that redness is an intrinsic property, such as surface spectral reflectance. But then the dispositional representationalist must say (as Shoemaker does) that the relevant contents attribute properties other than colors to objects. This is at least a counterintuitive thing to say.

We can sum up the problems for Russellian representationalism as follows. The following three claims are plausible: (i) color experiences attribute colors to objects; (ii) colors are intrinsic properties; (iii) there can be veridical phenomenally identical color experiences (in different subjects) of objects with relevantly different intrinsic properties. But these claims entail that phenomenally identical color experiences can attribute different properties to objects. So these claims are inconsistent with Russellian representationalism about color experience.

Denying any of these claims has significant costs. It is possible that some of these costs might be worth paying, if Russellian representationalism were the only way to get a viable representationalism off the ground. But I think that one can avoid these costs entirely. The problems arise only if we assume that the relevant contents must be Russellian contents. If we appeal to a quite different sort of content, the problems may be avoided.

7 Fregean content

Frege distinguished between the sense and the reference of linguistic expressions. The referent of an expression such as 'Hesperus' is the planet Venus, but its sense is a *mode of presentation* of that planet. Two distinct terms, such as 'Hesperus' and 'Phosphorus', may have the same referent but different senses. Something similar goes for a property term, such as 'bright'. Here one might say that the referent of the term is the property of brightness, while its sense is a mode of presentation of that property. Frege held that the sense of an expression determines the expression's referent, in some fashion.

One can make a similar distinction for mental states. A belief such as *Hesperus is bright* is composed of concepts, such as *Hesperus* and *bright*. The concepts can be said to have extensions: the planet Venus and the property of brightness, respectively. The concepts can

also be said to have *modes of presentation* of those extensions. These modes of presentation can be thought of as an aspect of the content of the concepts. Two different concepts, such as *Hesperus* and *Phosphorus* may have the same extension but different modes of presentation. The belief as a whole may have a complex content that is composed of these modes of presentation.

What is a mode of presentation? Frege was not entirely clear on this. But one natural approach characterizes a mode of presentation as a *condition on extension*. The idea is that every concept is associated with some *condition*, such that an entity in the world must satisfy this condition to qualify as the extension of the concept. For example, the concept *Hesperus* might be associated with a condition something like: *the object usually visible at a certain point in the evening sky*. (Strictly speaking, the condition is *being the object...*, but I will usually use the simpler form.) In the actual world, Venus satisfies this condition, so Venus is the extension of *Hesperus*. The concept *Phosphorus* might be associated with a different condition (involving the morning sky). Venus also satisfies this condition, so Venus is the extension of *Phosphorus*. On this analysis, *Hesperus* and *Phosphorus* have the same extensions but different modes of presentation, just as desired.

We can think of these modes of presentation as the *Fregean contents* of concepts. We can also speak of the Fregean contents of entire thoughts, by composing the Fregean contents of the concepts involved. For example, the Fregean content of the thought *Hesperus is Phosphorus* might be something like the condition *the object usually visible at a certain point in the evening sky is usually visible at a certain point in the morning sky*. This is a condition on the state of the world, reflecting what it takes for the thought to be true: that is, it is a sort of truth-condition.

To determine the condition of extension associated with a concept or with a thought, one can consider different hypothetical possibilities concerning the way one's world might turn out to be, and determine what the concept's extension (or the thought's truth-value) would be under those possibilities. For example, if it turned out that star A was the bright object in the evening sky and a different star B was the bright object in the morning sky, we would say that *Hesperus* picks out A and *Phosphorus* picks out B, and we would say that *Hesperus is Phosphorus* is false. Formally, one can say that the Fregean content of a concept is a mapping from scenarios to extensions, and that the Fregean content of a thought is a mapping from scenarios to truth-values, where scenarios are maximal epistemic possibilities, or centered possible worlds. I have developed this sort of approach extensively elsewhere (e.g. Chalmers 2002), but for the most part the informal understanding will suffice here.

(Note that modes of presentation are quite distinct from the manners of representation invoked earlier. Manners of representation are psychological features of an individual. By contrast, modes of presentation are a sort of *content*. Unlike manners of representation, modes of presentation have built-in conditions of satisfaction. Rather than being psychological entities, they are abstract entities to which psychological states may be related, by having them as their content.)

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

What are the modes of presentation associated with a given perceptual experience? To determine these, one considers scenarios involving different ways the world might turn out, and consider what the objects and properties represented by the experience will then be. Take a visual experience as of a green sphere. At a very rough first approximation, one might say that for an object to be represented by an experience, it must cause the experience in the appropriate way. So the mode of presentation of the object will be something like: *the object that is causing this experience in the appropriate way*. Likewise, one might say that for a property (say, greenness) to be attributed by the experience, it must be the property that has usually caused that sort of color experience in normal conditions in the past. So the mode of presentation of the property will be something like: *the property that usually causes phenomenally green experiences in normal conditions*.¹⁹

Overall (if we abstract from all other features of the experience, such as its spatial features), the experience will be associated with a Fregean content along the lines: *the object causing this experience has the property that usually causes experiences of phenomenal greenness*. This is very roughly compatible with our intuitive judgments about cases. If it

¹⁹ The suggestion that color experiences have dispositional modes of presentation of this sort that fix reference to intrinsic properties in the external world can be seen as analogous to a corresponding claim about color expressions: e.g., that “red” refers to an intrinsic property, with a mode of presentation that picks out whatever normally caused phenomenally red experiences.

turns out that no object is causing the experience in the appropriate way (e.g. if we are hallucinating the object), then we will judge that the experience is not veridical. If it turns out that the object has a property that normally causes phenomenally red experiences, but is now causing a phenomenally green experience due to abnormal conditions, we will judge that the experience is not veridical. And so on.

(One could also make a case for dropping the “object causing this experience” aspect of the content, yielding a “pure existential” Fregean content along the lines of “there exists an object [at such-and-such location] with the property that usually causes experiences of phenomenal greenness”. This turns on subtle issues about whether alleged cases of “veridical hallucination” (e.g. when one hallucinates a red square in front of one, and by coincidence there is a red square in front of one) should be classified as veridical or non-veridical. If veridical, one should go with the pure existential Fregean content. If non-veridical, one should go with the more complex causal Fregean content.)

One can formalize this approach to Fregean content by using the two-dimensional framework for analysing mental and linguistic content. I develop this sort of analysis in forthcoming work.²⁰ For present purposes, the informal understanding above will suffice.

8 Fregean representationalism

This leads to a natural proposal. Let us say that a *Fregean representational property* is the property of having a certain Fregean content (in a certain way). Let us say that *Fregean representationalism* is the thesis that phenomenal properties are equivalent to certain (pure or impure) Fregean representational properties. Then one might use the Fregean contents of the sort described above to put forward a version of Fregean representationalism.²¹

For example, one might propose that phenomenal redness is equivalent to the property of having a certain Fregean content (in the appropriate phenomenal way), where this Fregean content involves a mode of presentation such as *the property that normally causes experiences of phenomenal redness*. On this view, the relevant representational content does not directly involve the property attributed by the experience. It may well be that the experience attributes the property of redness to an object, and that redness is a surface spectral

²⁰ On this approach, the Fregean content of an experience can be analyzed as a primary or epistemic intension mapping epistemic possibilities to extensions. I develop this analysis in forthcoming work. In another forthcoming work, I try to develop a more refined view of the Fregean content of experience in a way that respects some of the phenomenological insight of primitivism.

²¹ Thompson 2003 also defends Fregean representationalism.

reflectance property. This attributed property may enter into the Russellian content of the experience, but it does not enter into the Fregean content. Rather, the Fregean content involves a mode of presentation of this property.

Here it is highly plausible that two phenomenally identical experiences will have the same Fregean content. This will certainly be the case if the Fregean content of perceptual experiences works as above, by requiring a certain relation to the relevant phenomenal property. For example, any two phenomenally red experiences will involve the mode of presentation *the property that normally causes phenomenally red experiences*. Of course, this mode of presentation may pick out different properties in different environments, or in creatures with different perceptual systems. In me, it might pick out a certain surface spectral reflectance; in a subject in a different environment or with a different perceptual system, it might pick out a different property entirely. But this is just what one would expect.

This suggests that a phenomenal property entails the corresponding Fregean representational property. For example, phenomenal redness entails visually phenomenally attributing a property under the Fregean mode of presentation *the property that normally causes phenomenally red experiences*. As for the converse, it seems plausible that any mental state that visually phenomenally attributes a property under this mode of presentation will itself be phenomenally red. One might be able to replace “visually” with “perceptually”, as it is implausible that a non-visual perceptual experience could attribute a property under this mode of presentation. It may well be that a *belief* could attribute a property under this mode of presentation, however, so one needs to at least specify a perceptual manner of presentation. As a result, it seems plausible that phenomenal redness is equivalent to the representational property of perceptually phenomenally representing the relevant Fregean content.

This view handles the central problems for Russellian representationalism straightforwardly. These views arose because Russellian representationalism is inconsistent with the plausible claims that (i) color experiences attribute colors to objects, (ii) color is an intrinsic property, and (iii) there can be veridical phenomenally identical experiences (in different subjects) of objects with relevantly different intrinsic properties. By contrast, Fregean representationalism is clearly compatible with these claims, for the reasons given above.

The view is plausibly a sort of narrow representationalism. I have argued elsewhere that the Fregean content of thoughts (understood appropriately) is always a sort of narrow content. It can equally be argued that the Fregean content of perception is a sort of narrow content. Given that the Fregean contents of a subject’s experiences are determined by the subject’s phenomenal properties, and that the phenomenal properties are intrinsic properties of a

subject, this conclusion follows automatically. So as long as the relevant manners of representation are also narrow, the resulting Fregean representational properties will be narrow, and we will have a sort of narrow representationalism.

The view is also naturally seen as a sort of nonreductive representationalism, as phenomenal notions are used in specifying the relevant Fregean contents. If one holds that the relevant manners of representation are also phenomenally specified (as I do), the result is a *doubly nonreductive* representationalism, involving phenomenal elements in both the manner of representation and in the content.

It seems that this view has the potential to accommodate most of the problems mentioned so far. It is compatible with the narrowness of phenomenology, while also being compatible with the view that experience attributes intrinsic properties to objects. It is also compatible with our intuitions about the veridicality or nonveridicality of experiences in specific cases. Fregean content is defined in such a way that most of these intuitions are accommodated automatically. If one judges that a certain experience in a certain environment will be nonveridical, this requires only that the Fregean content of the experience should be a condition that maps this environment to “false”. The overall Fregean content of the experience will be a condition that mirrors our judgments about the veridicality of an experience under arbitrary hypotheses about the state of the external world.

Of course, this means that the Fregean content may need to be refined beyond the very crude first characterization given above (involving *the property that normally causes phenomenally red experiences*). But this was only a rough approximation: the real content will be a condition (or mapping) that reflects our specific judgments about cases. So if there are cases that the crude characterization of the content misclassifies as veridical or nonveridical (as there certainly will be), one may want to refine the characterization, but in any case the problem concerns only the characterization, not the content. The general thesis that phenomenal properties are Fregean representational properties will be unthreatened by any such case.

This Fregean representationalism is closely related to Shoemaker’s dispositional Russellian representationalism, discussed earlier. Both views give a key role to dispositional notions such as *normally causes phenomenally red experiences*. The difference is that where Shoemaker’s view holds that color experiences attribute these dispositions to objects, my view holds that color experience attributes intrinsic properties (colors) to objects, with the dispositional notions serving as a mode of presentation of these intrinsic properties. That is, a dispositional relation to experience is used to *determine* the property attributed by the experience, but the property attributed is not itself dispositional. This avoids one major

difficulty for Shoemaker's view: the claim that the primary properties attributed by color experiences are not colors.

The other major difficulty for Shoemaker's view was the problem of specifying the relevant dispositions. This problem is easily handled here. The mode of presentation will be something like *the property that normally causes phenomenally red experiences in me*, where the last part involves an indexical mode of presentation of oneself, one that different individuals can use to refer to themselves.²² Indexical modes of presentation are already required for many other purposes in this sort of account of Fregean content, for example in specifying the Fregean content of indexical concepts such as *I* and *here*. Formally, they can be modeled using functions from centered possible worlds (worlds marked with an individual and a time) to extensions. So the Fregean content of *I* can be seen as a mapping that takes a centered world to the individual at the center of that world; and the Fregean content of a phenomenally red experience can be seen as a mapping that takes a centered world to whatever property normally causes red experiences in the individual at the center. While the notion of an indexical *property* is problematic, there is no such problem with the notion of an indexical mode of presentation. When a perceptual experience (embedded in an environment) has this sort of indexical mode of presentation, it will usually determine a perfectly objective property as its extension.

One might wonder whether this view is compatible with the oft-noted *transparency* of experiences. As I construe it, the central datum of transparency is that when we attempt to introspect the qualities of our experiences (e.g. phenomenal redness), we do so by attending to the qualities of external objects (e.g. redness). In effect, we look "through" the phenomenal property. But this is just one what one would expect where modes of presentation are involved. When one introspects the content of a belief such as *Hesperus is bright*, one does so by thinking about Hesperus; one looks right through the mode of presentation. But nevertheless the mode of presentation exists, and one can become introspectively aware of it. The same goes for manners of representation, such as visualness. One looks "through" this manner of representation on introspecting an experience of phenomenal greenness, but the manner is nevertheless there, and one can obviously become introspectively aware of it. (This case shows especially clearly that there is no inference from the datum of transparency to the absence of introspective awareness). So if phenomenal properties are associated with having

²² One might be tempted to replace *normally causes phenomenally red experiences in me* with *normally causes phenomenally red experiences in my community*. But this would have the odd consequence that someone who is spectrally inverted with respect to their community will have mostly illusory experiences.

certain modes of presentation of represented properties, under certain manners of representation, it is to be expected that these properties will be transparent.

One might object that the sort of modes of presentation I have been discussing over-intellectualize the contents of experience. When one attends to a red ball, one does not usually conceive of it as the cause of one's experience, or as possessing properties that normally cause that sort of experience. But this point is compatible with what I have said. The Fregean contents I have appealed to may very often be *nonconceptual* contents: to have a state with these contents, a subject need not deploy a concept with those contents. So a subject's visual experience can have a mode of presentation along the lines of *the object causing this experience* without the subject deploying the concept of causation, or of the experience. It should be kept in mind that modes of presentation are fundamentally *conditions*, which can be seen as mappings from scenarios to extensions. The descriptive characterizations such as *the object causing this experience* are only rough ways of characterizing these conditions. Of course we have to use linguistic expressions and concepts to formulate these characterizations, but the conditions themselves may be entirely nonlinguistic and nonconceptual.

Further, it seems clear that there is some psychological reality to these modes of presentation. When a subject has a perceptual experience, the subject is usually capable of judging whether the experience is veridical or not, depending on further information about the state of the world. For example, if a rational subject discovers that there is no object causing the experience, the subject will conclude that the experience is nonveridical. Speaking broadly, one might think of this pattern of judgments as part of the *inferential role* of the experience. The Fregean content of an experience reflects this sort of inferential role directly. And it is plausible that the (rational) inferential role of an experience is grounded in some fashion in the representational content of the experience. So there is good reason to think that Fregean representational properties play an important psychological role.

I have outlined how this account works for the case of a simple color experience. I think the account can be generalized to other visual experiences, though some tricky issues arise. Thompson (2003) argues that one can handle spatial experiences in a similar way: these function by attributing properties (or relations) to objects, where these properties are picked out under a mode of presentation that characterized them as being those properties that stand in the relevant causal relation to spatial phenomenal properties. Other aspects of visual experience might be handled in a similar way.

It is natural to extend the account to other perceptual experiences, and to bodily sensations. It is easy enough to see how this will go: in every case, there will be a mode of

presentation that picks out a property as something like the normal cause of the sort of experience in questions. Here, I will not try to address the issue of whether this extension would be appropriate. There are some difficult questions, not least about whether (for example) olfactory experiences or pain experiences can be veridical or nonveridical. I am inclined to say that *if* these experiences can be veridical or nonveridical, then an extension of the current approach will give a reasonable analysis of their content.

9 Summary

It may be useful to sum up the various distinctions that have been introduced, and to indicate where the view I have advocated falls.

Representationalism holds that phenomenal properties are equivalent to pure or impure representational properties: properties of representing a certain content (in a certain manner).

Pure representationalism holds that phenomenal properties are equivalent to pure representational properties: properties of representing a certain content. *Impure representationalism* holds that phenomenal properties are equivalent to impure representational properties: properties of representing a certain content in a certain manner.

Reductive representationalism holds that phenomenal properties are equivalent to representational properties that can be fully characterized in nonphenomenal terms. *Nonreductive representationalism* holds that phenomenal properties are equivalent to representational properties that cannot be fully characterized in nonphenomenal terms. Nonreductive representationalism might be nonreductive about the manner of representation, or about the representational content, or both.

Narrow representationalism holds that phenomenal properties are equivalent to narrow representational properties, depending only on a subject's intrinsic properties. *Wide representationalism* holds that phenomenal properties are equivalent to wide representational properties, depending partly on the subject's environment.

Russellian representationalism holds that phenomenal properties are equivalent to Russellian representational properties, with contents involving objects and/or properties in the world. (One can further distinguish *physical*, *projectivist*, *primitivist*, and *dispositionalist* Russellian representationalism, depending on what properties are held to be represented.) *Fregean representationalism* holds that phenomenal properties are equivalent to Fregean representational properties, with contents involving modes of presentation of objects and properties in the world.

The view I have advocated is a form of impure, nonreductive, narrow, Fregean representationalism. The view is in fact doubly nonreductive, with phenomenal elements involved in both the manner of representation and the representational content. In the case of a phenomenal color property, the associated manner of representation is required to be a perceptual phenomenal manner of representation. The associated representational content is a Fregean mode of presentation, one that picks out whatever property normally causes that sort of color experience in the subject. This view contrasts sharply with the most common versions of representationalism, which are reductive, wide, and Russellian.

10 Conclusion

I expect that the interface between consciousness and intentionality will be the central topic in the next decade of the philosophy of mind. I hope that the analysis I have given here helps to clarify some crucial issues in exploring this interface, and to open up some underexplored possibilities for making progress.

What of the issues mentioned at the start of the paper: is consciousness grounded in intentionality, or is intentionality grounded in consciousness? I have argued for a necessary equivalence between phenomenal and representational properties; but which of the phenomenal and the representational is more fundamental?

The nonreductive approach I have taken offers little prospect for grounding consciousness wholly in intentionality, where the latter is construed independently of consciousness. The representational analyses I have given all make unreduced appeal to phenomenal notions, and I think there is little hope of giving analyses that do without such notions altogether. So I think that a reduction of the phenomenal to the representational is not on the cards.

One might think that this approach offers more hope of grounding intentionality in consciousness. It is not implausible that there is something about consciousness, which by its very nature yields representation of the world. One might hold that at least with perceptual experiences, representational content accrues *in virtue* of the phenomenology. One might further hold that something similar holds for beliefs: their representational content accrues either in virtue of their phenomenal character, or in virtue of their connections to other beliefs and experiences whose content is grounded in phenomenal character.

Still, for this approach to provide a reductive grounding of the intentional, we would need to characterize the underlying phenomenal domain in non-intentional terms. And it is far from clear that this is possible. On the face of things, a characterization of my

phenomenology that avoids intentional notions entirely would be quite inadequate. Rather, intentional content appears to be part and parcel of phenomenology: it is part of the essential nature of phenomenology that it is directed outward at a world. If so, we cannot reduce intentionality to something more fundamental; at best, we can locate its roots in the intentionality of the phenomenal.

I think, then, that the most attractive view is one on which neither consciousness nor intentionality is more fundamental than the other. Rather, consciousness and intentionality are intertwined, all the way down to the ground.

Bibliography

- Block, N. 1978. Troubles with functionalism. In (C.W. Savage, ed.) *Perception and Cognition: Issues in the Foundation of Psychology*. Minneapolis: University of Minnesota Press.
- Block, N. 1990. Inverted earth. *Philosophical Perspectives* 4:53-79.
- Boghossian, P. & Velleman, J.D. 1989. Color as a secondary quality. *Mind* 98:81-103.
- Burge, T. 1979. Individualism and the mental. *Midwest Studies in Philosophy* 4:73-122.
- Byrne, A. 2001. Intentionalism defended. *Philosophical Review* 110:199-240.
- Carruthers, P. 2000. *Phenomenal Consciousness: A Naturalistic Theory*. Cambridge: Cambridge University Press.
- Chalmers, D.J. 1996. *The Conscious Mind: In Search of a Fundamental Theory*. Oxford University Press.
- Chalmers, D. (forthcoming). Perception and the Fall from Eden. In (T. Gendler and J. Hawthorne, eds) *Perceptual Experience*. Oxford: Oxford University Press.
- Chalmers, D. 2002. The components of content. In (D. Chalmers, ed) *Philosophy of Mind: Classical and Contemporary Readings*. Oxford University Press. Clark, A. 2000. *A Theory of Sentience*. Oxford: Oxford University Press. Crane 2003
- Crane, T. 2003. The intentional structure of consciousness. In (Q. Smith & A. Jokic, eds) *Consciousness: New Philosophical Perspectives*. Oxford: Oxford University Press.
- Dennett, D.C. 1968. *Content and Consciousness*. London: Routledge.
- Droege, P. 2003. *Caging the Beast: A Theory of Sensory Consciousness*. Amsterdam: John Benjamins.
- Dretske, F. 1995. *Naturalizing the Mind*. Cambridge, Mass.: MIT Press.
- Egan, A. (forthcoming). Appearance properties? <http://www.geocities.com/eganamit/APs.pdf>
- Harman, G. 1990. The intrinsic quality of experience. *Philosophical Perspectives*. 4:31-52
- Holman, E. 2002. Color eliminativism and color experience. *Pacific Philosophical Quarterly*. 83:38-56.
- Horgan, T. & Tienson, J. 2002. The intentionality of phenomenology and the phenomenology of intentionality. In (D. Chalmers, ed) *Philosophy of Mind: Classical and Contemporary Readings*. Oxford: Oxford University Press.

- Jackson, F. 2003. Mind and illusion. In (A. O'Hear, ed) *Minds and Persons*. Cambridge: Cambridge University Press.
- Kriegel, U. 2002a. PANIC theory and the prospects for a representational theory of phenomenal consciousness. *Philosophical Psychology*. 15:55-64.
- Kriegel 2002b. Phenomenal content. *Erkenntnis*. 57:175-98
- Levine, J. 2003. Experience and representation. In (Q. Smith & A. Jolic, eds) *Consciousness: New Philosophical Perspectives*. Oxford: Oxford University Press.
- Lloyd, D. 1991. Leaping to conclusions: connectionism, consciousness, and the computational mind. In (T. Horgan & J. Tienson, eds) *Connectionism and the Philosophy of Mind*. London: Kluwer.
- Loar, B. 2003. Phenomenal intentionality as the basis of mental content. In (Hahn & B. Ramberg, eds) *Reflections and Replies*. Cambridge, Mass.: MIT Press.
- Lycan, W.G. 1996. *Consciousness and Experience*. Cambridge, Mass.: MIT Press.
- Lycan, W.G. 2001. The case for phenomenal externalism. *Philosophical Perspectives* 15:17-35.
- Maud, J.B. 1995. *Colours: Their Nature and Representation*. Cambridge: Cambridge University Press.
- McGinn, C. 1988. Consciousness and content. *Proceedings of the British Academy* 74:219-39. Reprinted in *The Problem of Consciousness* (Oxford: Blackwell, 1991).
- McLaughlin, B. 2003. Color, consciousness, and color consciousness. In (Q. Smith & A. Jolic, eds) *Consciousness: New Philosophical Essays*. Oxford: Oxford University Press.
- Peacocke, C. 1983. *Sense and Content: Experience, Thought, and their Relations*. Oxford: Oxford University Press.
- Putnam, H. 1975. The meaning of 'meaning'. In *Mind, Language, and Reality*. Cambridge: Cambridge University Press.
- Rey, G. 1998. A narrow representationalist account of qualitative experience. *Philosophical Perspectives* 12:435-58.
- Rosenthal, D.M. 1997. A theory of consciousness. In (N. Block, O. Flanagan, and G. Güzeldere, eds) *The Nature of Consciousness*. Cambridge, Mass.: MIT Press.
- Searle, J.R. 1980. Minds, brains and programs. *Behavioral and Brain Sciences* 3:417-24.
- Searle, J.R. 1983. *Intentionality*. Cambridge University Press.
- Searle, J.R. 1990. Consciousness, explanatory inversion and cognitive science. *Behavioral and Brain Sciences*. 13:585-642.
- Shoemaker, S. 1990. Qualities and qualia: What's in the mind? *Philosophy and Phenomenological Research Supplement*. 50:109-131.
- Shoemaker, S. 1994. Phenomenal character. *Nous* 28:21-38.
- Shoemaker, S. 2001. Introspection and phenomenal character. *Philosophical Topics*. Reprinted in (D. Chalmers, ed) *Philosophy of Mind: Classical and Contemporary Readings*. Oxford: Oxford University Press.
- Siewert, C. 1998. *The Significance of Consciousness*. Princeton: Princeton University Press.
- Stoljar, D. (forthcoming) Consequences of intentionalism.
<http://philrsss.anu.edu.au/~dstoljar/onlinepapers/Consequences.pdf>

- Strawson, G. 1994. *Mental Reality*. Cambridge, Mass.: MIT Press.
- Thau, M. 2002. *Consciousness and Cognition*. Oxford: Oxford University Press.
- Thompson, B. 2003. *The Nature of Phenomenal Content*. Ph.D. dissertation, University of Arizona.
- Travis, C. 2004. The silence of the senses. *Mind*. 113:57-94
- Tye, M. 1995. *Ten Problems of Consciousness: A Representational Theory of the Phenomenal Mind*. Cambridge, Mass.: MIT Press.
- Tye, M. 2003. Blurry images, double vision, and other oddities: New problems for representationalism? In (Q. Smith & A. Jokic, eds) *Consciousness: New Philosophical Perspectives*. Oxford: Oxford University Press.
- Vinueza, A. 2000. Sensations and the language of thought. *Philosophical Psychology*. 13:373-392.
- Warfield, T. 1999. Against representational theories of consciousness. *Journal of Consciousness Studies*. 6:66-69.
- Wright, W. 2003. Projectivist representationalism and color. *Philosophical Psychology*. 16:515-533