

Saving Appearances: a Dilemma for Physicalists

Charles Siewert
University of California, Riverside
March 30, 2008

(To appear in *The Waning of Materialism*, edited by G. Bealer and R. Koons,
Oxford University Press)

1. Introduction

The sort of physicalism or materialism I will criticize finds expression in slogans of the form: 'X is *nothing but*—or *nothing over and above*—Y.' What goes in for 'X' is, for example, consciousness, or some type of experience—at any rate, it is something mental or psychological. Such views are plausibly described as "reductionist." By purporting to tell us what every mentalistic X *is*, they promise to *explain* it in physical terms. But just what makes something physical in the relevant sense is not obvious, and partly for this reason it is hard to see how precisely to characterize these views as a group. However, that does not matter too much for my purposes, as long as the difficulty I will spell out does in fact arise for views I will describe, whatever general category might unite them.

I have advertised the issues broadly (using 'consciousness,' 'mental,' and 'experience'), but my discussion will not remain this general. Rather I will concentrate more specifically on the notion of *visual appearance*. I will describe scenarios contrasting the occurrence of visual appearance in a certain sense with its absence. This will provide us with both a way to clarify our understanding of visual appearance and a test for theoretical recognition of its reality (and, by implication, of the reality of phenomenal consciousness generally). This then will motivate a challenge to "ambitious" or reductive physicalist/materialist theories. Roughly speaking, the problem will be that, depending on the form they take, they either ultimately run into trouble acknowledging the reality of appearances, or else, to save appearances while holding onto their ambitions, they purport to discover deep necessities where warrant for these eludes us. If this dilemma is not resolved, then we should abandon ambitions to reduce consciousness to something physical. The question that finally confronts us then is whether this gives us any cause for regret.

2. 'Looks' in a sensory, phenomenal sense

To understand the challenge I want to pose, we must put off discussing the theories in question for awhile, and first try patiently to clarify the relevant notion of visual appearance. By 'a visual appearance' I mean an instance of something's *looking somehow to someone*. The use of 'looks' I want to invoke can be found in sentences of the form 'x looks F (e.g., green, or oval) to S.' But it is not confined to these. It can be found when I make claims like 'It looks to me

as if this figure is larger than that one,' or: 'That one looks to me changing in shape, and *this* one looks to me to be staying the same shape,' or: 'That looks to me like a hat.'¹ When thinking of such examples, note that it is a *sensory* (visual) sense of 'looks' to which I wish to make appeal. Not all "looks" are sensory. There is a (non-sensory, non-visual) sense in which we can speak of how it looks to us even when we have our eyes closed in a dark room, or if we are totally blind—as when someone says, 'It looks to me as if the next election's going to be very close.'²

Now I want to focus on a sense of 'looks' and 'visually appear' that is not only *sensory*, but *phenomenal*. To clarify this, I will not contrast a phenomenal sense of 'looks' with a non-phenomenal one, but with a non-phenomenal sense of 'vision.' One can be said to have *visual perception* or form *visual representations*, or have *vision*, even in the *absence* of visual appearances. There is a sense in which you might be said to *see* things, even when they do not *look* any way to you. Let me try to make clear how this could be by reference to the phenomenon of "blindsight." Psychologists and philosophers have extensively discussed experimental work on "blindsighters"—subjects who have retained some visual functioning while suffering certain kinds of damage to their visual cortex. But not everyone describes such subjects in just the same way. I want to make explicit a way of understanding the condition of blindsight subjects that will help illuminate the notion of visual appearance at issue. This will take a little while to set out fully.³

We start with this. To consider the condition of a blindsighter, suppose you are seated before a screen, in a well-lit room with your eyes open, and in various locations on the screen various relatively simple figures and arrangements of lines are briefly presented—such as X's and O's, or rows of either horizontal, or vertical lines. Now suppose that because of damage to your brain, a change has occurred in your reports about what you see when such stimuli are presented. On some portion of the screen (on the left, let's say), where, with gaze fixed, you formerly would have reported seeing an X or an O (or horizontal or vertical lines), now, post-trauma, you no longer report seeing anything at all. When an X is flashed there, and you are asked what you see there, you say 'Nothing.' Nonetheless, when asked which of two figures (X or O) was presented there, and *required to make a choice*, you turn out to answer correctly at levels significantly above chance. Somehow the effect of the light from the stimulus on your eye and brain has enabled you to make this sort of "forced" or "prompted" discrimination of the presented figures. There is, it seems, a sense in which we can regard this as a case of "visual" discrimination. This is blindsight. That such cases are actual has been made clear by the research of Weiskrantz (1986, 1997) and others.

You might understand a bit more specifically just what this condition involves in the following way. Suppose that, in the situation just described, the stimuli in that area where you deny seeing anything do not, in fact, *look* or *visually appear* any way at all to you—in a sense in which, elsewhere in the area

before you, things still *do* look to you somehow. If you were to say that *nothing in that region of the screen looked any way at all to you*, you would be speaking the *truth*. And so, provided that talk of ‘seeing’ is interpreted so as to entail that you can’t see what doesn’t *look* any way to you, your denial that you see anything there would also be correct. Still, your brain’s visual system enables you to make forced choice discriminations of the shapes or orientations of these figures, and so perhaps we should say there is also a sense in which you “see” or “visually perceive” them. Thus: in one sense you are *blind* to the figures (they don’t look any way to you); in another sense you *see* them (you somehow discriminate and perhaps represent them by means of your brain’s visual system)—hence you have “blindsight.” At any rate, this is an intelligible way of conceiving of the condition of blindsight.⁴

This shows we can conceive of states that are in *some* sense visual (and thus sensory) that are, however, not *phenomenal* visual states—not visual *appearances*. One may suppose that there is some *visual* state (some state of the visual system)—even one we may regard as “representational” or “information-bearing”—that is present in blindsight, and enables one to make, when prompted, discriminations of the figures one denies seeing. To grasp the contrast between vision of this sort, and that in which something looks some way to you (in the sensory sense) is to take a crucial step in grasping what I mean when I say visual appearance is phenomenal.

Now let’s refine our understanding of visual appearance further, by altering the features of the blindsight situation a little, so as to consider a *hypothetical* form of it. Note that discriminations made in the sort of blindsight just discussed are elicited by means of “forced choices”—prompting with a choice ‘X or O?’ offered to subjects who regard their responses as mere “guesses.” However, it is not very difficult to conceive of a form of blindsight in which the discriminations, as revealed in verbal reports, button-pressings and the like, arise *spontaneously and unprompted by others’ multiple choice questions*. Suppose that—while it remains true that the stimuli in your upper left quadrant do not look or visually appear any way to you—still, you find yourself disposed to judge, without having a choice forced on you, that a figure of a certain shape or orientation is or was just there—for some (perhaps not very extensive) range of shapes, orientations, and locations. For example, it just spontaneously occurs to you (perhaps in response to a self-posed question, perhaps not) that there is an X there, say. And you feel confident that this is true, as you might feel confident in your spontaneous and direct answer to the question, ‘What day of the week is today?’ And we may suppose: you are inclined to point to where you think the stimulus is, and adjust your gaze so as to look *at* it (in which case it *does* then look X-shaped to you, and you thereby confirm your judgment). So you also have this *non-verbal* responsiveness to the stimulus. Thus, to interpret ‘looks’ in the phenomenal manner, I ask you to conceive of this sort of spontaneous visual response to a stimulus that does *not* look any way to you, and contrast this with a case in which it *does*.

But here please note: I am *not* asking you to assume that spontaneous blindsight judgment is compatible with the principles governing vision (e.g., human neuropsychology) in the actual world. Maybe it is not—and the case is “purely hypothetical.” However, I do assume that for the purposes of clarifying how we understand an expression, we can consider how it would apply (or not apply) in at least some hypothetical cases, without needing to be confident that the situations described wouldn’t demand the natural world somehow be other than it actually is.

To get clearer about how to understand this hypothetical case, I need to make this explicit. I am not asking you to suppose that *what* is judged, the *content* of spontaneous judgments made in the two situations—with and without visual appearance—would be *in every way indistinguishable*. That might be impossible. To see this, consider: you may have no way of thinking specifically of the peculiar shape you see something to have on a given occasion, but demonstratively, as: *this shape*. And how you understand *which* shape *this* one is may depend on the stimulus right then and there *appearing* to you a certain way. You might articulate this way of understanding by saying that *this* shape is the shape that *this now looks* to you. Similarly we might say that you can judge a stimulus to have a certain *color*—*this shade*—where your understanding of what you mean by ‘this shade’ depends on its looking some way to you then and there: *the shade this now looks to me*. Thus I may understand which shape or shade I take a stimulus to have, in a way properly expressed by speaking of its concurrently *looking* somehow to me.

But it seems that your spontaneous blindsight discriminatory judgment could not have just the content identified relative to this way of understanding. For, *by hypothesis*, you are then *lacking* the visual appearance required. And plausibly, you cannot understand which shape you are thinking of, in a manner to be articulated as: *the shape this now looks to me*—if in fact *this now looks to you no way at all*. The point is: when we consider the scenario of spontaneous blindsight discrimination, we don’t assume that we can entirely *duplicate* in blindsight all the very same judgments about visual stimuli that phenomenal vision affords. For we may wish to recognize a class of “directly appearance dependent” judgments. These would be judgments where an expression of *what is judged* is understood by the subject in a manner that requires that it look or visually appear somehow to the subject at the very time of judgment. Subjects *lacking* the relevant appearances (like blindsighters) obviously could not make such judgments, or engage in behavior expressive of them.⁵

Clearly though, there will be *some* ways of classifying stimuli that can figure in spontaneous visually triggered judgments that are *not* “directly appearance-dependent” in the way just illustrated. You can, of course, somehow understand what you mean by an assertive utterance of ‘O-shaped,’ ‘vertical,’ ‘green,’ ‘brighter,’ ‘larger,’ or ‘above’—without its *then and there looking* to you in

ways thus reportable. And this remains so, even in the case where something looks O-shaped to you, and you judge it to be O-shaped: here too your understanding of what you mean by 'O-shaped' does not require that something then and there look O-shaped to you. For you understand what you mean by this term in the same way when you close your eyes. Your understanding of classifications (both in the blindsight and in the normal cases) can thus be at least *relatively* independent of current visual appearance, in the sense that they are not directly appearance-dependent.

You have so far supposed that you possess a power of spontaneous judgment to discriminate at least some features in part of your visual field where you would correctly report that nothing *looks* any way to you. Now however, to simplify further discussion, make the part of your visual field in which nothing looks any way to you encompass its *entire extent*. *Nothing at all* looks any way to you. It follows of course that the way you understand the terms by which you would express any spontaneous judgment about the stimuli—'X,' 'O,' 'vertical line,' 'horizontal line'—is *at least relatively independent* of visual appearance. (They are "at least relatively independent" if they aren't *directly dependent*, and in this case they can't be directly dependent on your enjoying visual appearances, because you *haven't got any*.)

Now—compare this kind of blindsight case with a rather seriously degraded sort of phenomenal vision. Suppose nothing in your visual field looks any way to you, *except* what occurs in this relatively small part, and what appears there looks very blurry, and your vision is severely lacking in acuity. Still you would say correctly that something *looks* some way to you—perhaps you can say whether some patch of light looks brighter or darker, larger or smaller, moving or stationary. But stimuli there do not look distinctly X-shaped, O-shaped, etc., and the lines whose orientation you are asked to judge are not distinctly apparent to you. There seems now to be no obstacle to supposing that the range of spontaneous *blindsight* judgments you can make, in the case where nothing looks anyhow to you, are not overall inferior to those one can make in the latter "blursight" case, in terms your understanding of which is not directly dependent on visual appearance. In fact, we can suppose that the judgments made with blindsight (e.g., of shape, size, and orientations) would be in some ways better (more discriminating) than the relatively appearance-independent judgments made with phenomenal blursight.⁶

So the contrast here is between a subject to whom nothing (sensorily) *looks* any way at all, and one to whom only stimuli in a relatively small area before her looks any way at all—and then only very blurrily, and in a manner that affords only relatively crude discriminatory judgments. One can conceive of the former subject (let's now call her "Linda") having powers of blindsight that enable her to make unforced discriminatory judgments that are in some ways better than the relatively appearance-independent judgments made by the latter (call her 'Glenda').

Now for one more twist. We have recognized interpretations of ‘visual,’ ‘vision,’ and ‘sight’ that allow these to apply where there is discrimination of stimuli wherein what is “seen” does not *look* or *visually appear* to the sighted person. But if we can make such a notion of “the visual” available to ourselves, there is no reason we cannot make it available to the subjects in our scenarios. Suppose now that Linda can also classify the judgments she makes as “visual” in the sense we have just recognized. Much as, seemingly without inference, she is able to tell when she is *judging* that something is so, as opposed to *wondering whether* it is, she is also able to classify her blindsight judgments as *visual ones*. We might say then that Linda not only enjoys a spontaneous blindsight—her blindsight is “*reflective*” in this sense: she can make discriminatory judgments not only about the visual stimuli but about these very discriminatory judgments as well, and thereby classify them as visual.

Glenda, meanwhile, though she can say that things look somehow to her when they do, has no inclination to employ *this* concept of the *visual* in her judgments about herself. Instead she is disposed to say that things do or don’t *look* some way to her, and that they are thus and so, because of how they look to her, and that she judges that things look some way to her—where her use of ‘look’ is to be interpreted in line with our previous remarks. In this respect, I might add, Glenda and I are not so different. For my ordinary reflective first-person judgments about vision are about things *looking* some way to me, or about my *seeing* this or that (in a sense of ‘see’ that implies visual appearance). Although on reflection I can acknowledge a way of understanding ‘vision’ and ‘seeing’ divorced from looking, that is not a concept of “visual” I ordinarily employ in thinking about my own vision.

The immediate point of this exercise was to explain how to interpret ‘looks’ and ‘visual appearance’ in a certain way. If it was successful, you recognize a sensory, phenomenal sense in which things look to you somehow, which allows you to contrast the situation of Linda (to whom things don’t look any way) with that of Glenda (to whom things do, but only quite blurrily)—while nevertheless Linda is able to make unprompted discriminatory judgments of somewhat greater range than the “relatively appearance-independent” judgments Glenda can make of shape, size, etc. Of course, one may refuse to cooperate in this exercise. But if you show a little patience, I don’t think you need find any insuperable difficulty in interpreting ‘looks’ or ‘appearance’ in the manner I suggest.

3. A Test for Phenomenal Realism

In order to pose my problem for ambitious physicalism, I need to make a couple of points explicit. Recall that I have not been asking you to assume that the occurrence of two subjects such as Linda and Glenda is consistent with the principles governing the occurrence of visual perception in the actual world. Let us say that if their occurrence is *not* consistent with those principles, then the contrast described between Linda and Glenda is “nomologically impossible.” That

is, in that case it would be nomologically impossible for two subjects to differ as they were said to differ with respect to the visual appearances they enjoyed while otherwise being just as they were said to be.

Now, it seems that acceptance of certain claims about what we *mean*, what we *have in mind*, or what *concepts* we use, when we describe something in a certain way, would commit us to saying that certain scenarios are *inconceivable* for us in a sense that would have us regard them as *more than just nomologically impossible*. If I asked you to conceive of a situation in which Linda did not differ from Glenda in height, but Glenda was taller by an inch, you might reasonably respond that this is inconceivable for you. For when you say 'x is taller than y' part of what you mean or part of what you have in mind is that x and y *do* differ in height. Or it is part of your concept of *being taller than* that they would differ in height. In saying this, you would not be saying just that the occurrence of two such subjects is inconsistent with principles you accept governing the conditions under which objects in the actual world vary or are similar in size. You would not just be saying you regard the scenario as nomologically impossible.

Similarly, switching from talk of *meaning* and *concepts*, to talk of *natures*, *properties*, and *constitution*—if you said that differing in height belongs to the *property of being taller*, or is part of *what it is to be taller*, or partly *constitutes being taller*, then you would be saying something stronger than just that, as things are in the actual world, things that differ in tallness differ in height. You would be saying more than just that it is nomologically impossible to differ with respect to the one without differing with respect to the other. Now perhaps “the right sort of” principles of nature secure truths about the identity of properties, natures, or what constitutes what. But the sheer fact that there is a nomological relationship between F and G (even one that says you get F *just when* you get G) is not by itself sufficient to secure such truths. So it seems that there is more than a mere nomological necessity and possibility invoked when one speaks of the identities of properties, what the nature of something is, or what constitutes what. And if this cannot somehow be accounted for in terms of what is conceivable, or *conceptually* necessary and possible, we may wish to speak here, as one sometimes does, of “metaphysical” possibility—the sense in which it has been claimed that it is *metaphysically* impossible to have water without H₂O.

I rehearse these points because I am concerned that certain views about mind apparently (if not always explicitly) put themselves in the business of making either conceptual or metaphysical assumptions that would commit one to claims of more than nomological impossibility. Such theories would commit us to saying that the contrast I have invited us to consider between Linda and Glenda is more than just nomologically impossible—it's somehow rationally *inconceivable*, or perhaps *metaphysically* impossible. But that is not a conclusion I think we should accept lightly. For I believe that the way of interpreting 'looks' manifest in considering that contrast is not arbitrary, but provides a way of

articulating and refining an understanding of ordinary talk of things looking this or that way to us. But if we do understand 'looks' in this manner, and yet it turns out that this interpretation should be deemed somehow conceptually incoherent, or that Linda's form of blindsight not only can't happen, given actual world psychology, but in some more fundamental way *couldn't* happen ("in any possible world"), then it turns out that our belief that things look somehow or other to us is deeply confused—there is no difference, of the sort we suppose, between the presence and absence of visual appearance; we literally don't know what we're talking about, when we say things look somehow to us. If the sort of contrast we seem to draw between Linda and Glenda really should be regarded as inconceivable, or if there is just nothing in reality that will permit any such contrast, then visual appearance is a cognitive illusion, and phenomenal vision is not real.

So I would propose that a theory's conceptual and metaphysical tolerance for the contrast between Linda's blindsight and Glenda's impoverished phenomenal vision provides a test for recognition of the reality of phenomenal vision itself. If acceptance of a theory would deprive you of the ability even to conceive coherently of this contrast, or it would commit you to regarding this contrast as deeply (more than nomologically) impossible, then it implicitly denies the reality of phenomenal vision.

4. Physicalism Challenged

Now I want to explain why some views don't seem to me to pass this test. Consider some of the remarks made by Daniel Dennett. He warns us against the "trap" of thinking that, in vision, there is any appearance to be reckoned with—any "seeming"—over and above one's "judging" in some way that something is so. So when faced with a case of visual color illusion, in which the spaces between lines in a grid seem pink, he says it would be a mistake to think there is "a phenomenon of really seeming—over and above the phenomenon of judging in one way or another that something is the case." (Dennett, 1991, p. 364) The manner of judging in question is likened to a "presentiment": it just occurs to you "that there is something pink out there, the way it might suddenly occur to you that there's somebody standing behind you." So: if you think that, in vision, things look some way to you, where this is distinguishable from your somehow spontaneously judging them to be some way, then you're wrong. This is important for Dennett because much of the point of his book *Consciousness Explained* is to persuade the reader that there is in reality nothing of the sort that people confusedly suppose a theory of consciousness needs to explain. We shouldn't ask a theory of consciousness to explain "real seeming," since there *just isn't any such phenomenon*.

But it is not always so clear exactly what Dennett is denying with his "nothing over and above" talk. It seems he is not denying that we can conceive of someone spontaneously judging a visual stimulus is there, when it does not

visually seem any way to her. (He agrees that a form of spontaneous blindsight is conceivable.) Rather, it seems he's saying that the contrast between such judgment and a visual appearance is *really nothing but a contrast between two species of "judgment"*—differing in the fineness of their content. For what he says he finds inconceivable is that *what is judged*—the *contents* of a spontaneous blindsighter's judgments—might surpass the "*limited and crude*" sort found in actual blindsight. To help motivate this, he suggests that we consider the "talents" that actual blindsighters, with their crude and limited repertoire, lack with regard to the visual stimuli they deny seeing. Then he contends "that what people have in mind when they talk of 'visual consciousness,' 'actually seeing,' and the like, is nothing over and above some collection or other of these missing talents." (Dennett, 1993, p.151) So, it seems the claim is that "all we have in mind" (or perhaps: all we *should* have in mind) when we contrast visual appearance with its absence is this difference in talents that goes with a difference in the relative "fineness" or "crudity" of the contents of judgments.

But at this point recall: Glenda, to whom we supposed things do look some way, has a pretty limited and crude content repertoire. And the only contents we supposed she had in her repertoire that Linda lacked were those the understanding of which depended on things looking somehow to her, in a sense of 'looks' specified by way of contrast with spontaneous verbally expressible judgment. But the contents of Glenda's judgments were not generally *finer*, or *less crude*, than Linda's. In fact, in some respects, Linda had a greater range of contents available for judging. Glenda's judgments differed from Linda's not in *greater general refinement*. They differed in depending on stimuli *looking* somehow to her—because they were judgments understood by her in a manner she might express by saying, e.g., 'the size (or shape or brightness) *this looks to me*.' However, it seems as if Dennett is saying a difference in how talented one is with respect to visual stimuli, how finely or crudely one's judgment distinguishes them, is all we do or should have in mind when we talk about visual appearance and its absence. But if we accept this, our scenario becomes incoherent. We tried to conceive of a scenario in which one subject (Glenda) is *not* superior to another (Linda) in respect of certain talents, when really all we had (or should have had) in mind when we distinguished them is that the first is more talented than the second in those respects. Belief in Dennett's view of consciousness would in that case render inconceivable the contrast wherein our grasp of the notion of visual appearance is evident, and so implicitly deny the reality of phenomenal vision.

Similar points might be applied to the general strategy for thinking about the mind, sometimes explicitly used for defending a physicalist or materialist perspective, known as functionalism—for Dennett's view might be seen as a variant of this strategy. There is great diversity in what may be regarded as a functionalist theory of mind. But I take the basic idea to be that differences in what mental states you have amount to differences in what "functional roles" are being played or realized in you—where great variation is allowed in the physical

nature of the structures or mechanisms that can play or realize such roles. In the case of visual appearance, the proposal would be that the difference between something's looking somehow to you, and its looking some other way to you (or not at all), is accounted for entirely in terms of some "functional" difference, plus whatever physical mechanisms realize just those functions. The functionalist will say there is nothing to vision over and above exercising certain "multiply realizable" functions (i.e., they can be realized in a physically very diverse range of entities), which functions can be understood ultimately in non-phenomenal terms—without employing 'looks' or 'visual appearance' in the sense to be accounted for. Once you have explained what these functions are, and what physical mechanisms carry out these functions in a particular case (or a particular "population")—you have explained everything about visual appearance.

This suggests that functionalism will also have trouble passing the test for phenomenal realism. If there is to be really *wide* variation in the internal physical make-up of entities that realize the functions in question, it will seem that the character of the functions will have to be, in a sense, "manifest," not "hidden." That is, they will be the sort of functions one can have warrant for thinking are somehow realized in oneself or someone else, even without internal, beneath the surface examination of the entity in whom they are realized—as one could have warrant for regarding Linda as capable of crude spontaneous visual discrimination without directly or indirectly observing her brain. But now if the functions that constitute vision (or our concept of it) are entirely *manifest* (only their physically variable realization is hidden), then it will be hard to see how to avoid rendering the contrast between Linda and Glenda either inconceivable or more than nomologically impossible.

For consider: if functionalism maintains a strong claim of multiple realizability, then it would seem to need to hold at least one of two claims. The "conceptual" functionalist would hold that "all we really have in mind" (or "all we mean") by saying that things would look some way to Glenda but not to Linda, is that Glenda implements a certain manifest function (talent, disposition) absent in Linda. And a "manifest property" functionalist would hold that the subjects' difference in experience would *consist in nothing but* Glenda's implementing a manifest function, absent in Linda. The worry then is that we conceived of Linda and Glenda in a manner that seems to leave us no way to spell out what we mean by saying it looks somehow to one but not the other, or what constitutes this difference, in terms of some manifest functional superiority on the part of Glenda, *without resort to the phenomenal difference between them*. For once we see that Glenda's phenomenal vision is not functionally superior to Linda's blindsight, with respect to the general fineness of discrimination it affords, we will find that the only manifest differences that would cling to such subjects, *relevant* for differentiating them phenomenally, are those whose characterizations depend on thinking of its *looking* somehow to Glenda but not to Linda. No one would or should maintain, for instance, that all we *mean* by or what *constitutes* the difference between them is the difference in the phonological properties by which

they express the difference in appearance-dependent attitudes. That is, no one should hold that Glenda's *phenomenal* superiority lies in her being disposed to utter *types of sounds* the Linda is not disposed to utter. And we shouldn't think we can conjure a difference in whether things look somehow to the two subjects *purely* out of the fact that they merely *think* there is that very difference between them. We shouldn't think that what gives Glenda something more—visual appearances—than Linda, is simply that Glenda *thinks* she's got something more.⁷ So to locate a functional superiority, we are ultimately thrown back on the fact that Glenda has attitudes of a sort that can be had only by someone to whom something *looks* somehow. But to rely ineliminably on differences so characterized (i.e., phenomenally) would be to reduce the theory to triviality. The *functional* difference in terms of which the *phenomenal* difference was to be identified would then be itself identified partly in terms of the phenomenal difference in question. Thus to confine ourselves to the resources for thinking about the mind to which a *non-trivial* manifest functionalism would restrict us would be to render inconceivable or deeply impossible the contrast between Linda's and Glenda's vision. So, unless we can point to some manifest functional difference in the contrast that would supply a reasonable candidate for capturing the phenomenal difference in non-phenomenal terms, we should think that such functionalism does not pass the test for phenomenal realism.⁸

We can see perhaps more clearly how this sort of problem arises if we consider certain views that are rather more specific than the generic, programmatic functionalism just considered. I have in mind here accounts whose authors (or critics) tend to classify as "reductive representationalist" theories of phenomenal character. These appear to come in two general varieties: "first-order" representationalist, and "higher-order" representationalist. The first would hold that visual appearance—its looking somehow to someone—can be reduced to (and thereby *explained as*) a special form of mental representation of objects and features in one's environment ("first-order" representations). The second sort of view would hold that phenomenal vision requires (and is to be explained as) a certain way of representing one's own visual states (hence "higher-order" mental representations). These views' explanatory ambitions would bring them under the vague rubric of physicalism insofar as they promise to fit phenomenal sensory states into a general "naturalistic" theory of mental representation that accounts for this in physical terms. I will take as an example of "first-order" reductive representationalism the theory articulated and defended in some detail by Michael Tye (1995, 2002), while my "higher-order representationalist" will be David Rosenthal (2002), who has developed his "higher-order thought" version of the idea in detail.

According to Tye, the phenomenal character of sense perception is explained by reducing it to what he calls "poised, abstract, non-conceptual intentional content," which he refers to by the acronym "PANIC." The theory holds that something looks somehow colored and shaped to you—thus you have experience with a certain "phenomenal character"—just when your brain forms a

kind of fine-grained and unified representation of color and shape that represents in a “map-like” fashion (by contrast with the “sentence like” manner of *conceptual* representations), and which is “poised” to affect your beliefs (that is, available to have a direct impact on your beliefs by, e.g., disposing you to judge that there is something of a certain color and shape before you).

But now, if we adopt this theory, and think that its looking somehow to you is nothing but your being in a “PANIC” state, how will we conceive of the contrast between Linda’s blindsight and Glenda’s blursight? What has Glenda got that Linda doesn’t? Linda’s blindsight is no worse overall than Glenda’s poor phenomenal vision in what it affords by way of spontaneous discrimination of visual stimuli. So we would seem to have just as much call to think that Linda has encoded in her head a map-like representation of the visual stimuli to which she responds, which is not less detailed and unified than any we would be entitled to attribute to Glenda. And Tye himself would admit that there is nothing to prevent one from forming this sort of non-conceptual representation of what looks no way at all to one—otherwise there would be no need to introduce the “poise” condition for PANIC. So would they differ with respect to *poise*? Recall that our two subjects are both disposed to spontaneous judgments, verbally and otherwise expressible, about what types of stimuli are before them. And so there would seem to be nothing relevant to distinguish their “poise”—except that Glenda would be disposed to judgments about (e.g.) shape and color, her understanding of which would be articulated by talking about how things looked to her. But we don’t want to have the theory characterize the special sort of representational content that allows us reductively to explain things looking to us as they do by saying—“and it’s the kind of content that can be had only when things *look* to one somehow.” That kind of circularity would undermine the theory’s right to claim it *explains* visual appearance. So, if we are to distinguish the two subjects content-wise, it will need to be with some “phenomenally-purified” notion of their contents. But once we set aside the difference in their appearance-dependent judgments, the two subjects are not to be distinguished in terms of the “fineness of grain” of content.

What are we to conclude? Linda would have “non-conceptual representations” of what’s before her affecting eyes, poised to affect her beliefs, crude but generally no less fine in content than Glenda’s, except insofar as Glenda’s grasp of the content of her belief is conceived to depend on things *looking* somehow to her. To be sure, Glenda enjoys more finely grained *visual appearances* than Linda, and makes more finely grained judgments dependent on these than she does, since Linda does not enjoy visual appearances at all. But the adoption of first-order reductive representationalism about phenomenal vision would not allow us ineliminable appeal to a difference in how things look to the subjects, in order to conceive of the difference between them (on pain of abandoning the theory’s explanatory ambitions). Thus, its adoption would leave us no resources for supposing that things don’t look any way to Linda, but do to Glenda. Moreover, the theory would tell us that the two subjects, given their

visual talents, would have all they needed to be subjects of phenomenal vision—it would look somehow to both of them—by force of *metaphysical necessity*. There is no possible world in which things look blurry to Glenda on the left, but no way at all to Linda. Thus adoption of the theory would render the contrast between the two subjects inconceivable to us, and would declare the contrast itself a metaphysical impossibility. So the theory does not pass my test for phenomenal realism.

Now higher-order reductive representationists would have something to say about what goes wrong in Tye's sort of theory. They would say he fails to recognize the need for the subject to *represent her own visual state*. For *that* is what constitutes the difference between things looking blurry to her and their looking to her no way at all. Following Rosenthal's account, we would say something like this. Phenomenal vision consists in the union of two otherwise dissociable aspects: the state of having of *sensory (visual) qualities*, and the having of *seemingly non-inferential, assertoric thoughts* to the effect that one is in such a visual state ("higher-order thoughts"). Visual qualities (and sensory qualities generally) are understood to be qualities of sensory states, in virtue of which we can distinguish and classify them, instantiable in the absence of phenomenal vision (or phenomenal sense experience generally). For as Rosenthal conceives it, the condition that enables subjects to discriminate stimuli in blindsight (or in subliminal perception) has visual qualities. (Similarly, for him your state of being "in pain" has a sensory quality of some sort, even when (because of distraction) it doesn't *feel* any way to you to be in pain.) And so, since we're interpreting 'look' in such a way that things don't look any way to the blindsighter, possession of visual qualities is insufficient for having things look somehow to you. If we adopt the Rosenthal view, we need to say that what would make the blindsighter, possessed merely of visual qualities, someone to whom things look some way, would be the addition of seemingly non-inferential assertoric thoughts to the effect that she is in a state so visually "qualified."

Recall again the contrast we imagined between Linda and Glenda. Linda we supposed to be a *reflective* blindsighter: she spontaneously judges that her discriminations of stimuli are "visual"—in the sense of 'visual' we recognized when we said that there can be a sort of vision in the absence of visual appearance. Glenda, on the other hand, we imagined inclined to no judgments about her own visual condition other than those to which I said I was ordinarily inclined: we are both disposed to say things look to us somehow or other, or that we do or don't see them (in a sense in which their looking no way to us precludes our seeing them). But even if there is vision without visual appearance in this sense, that's not a kind of vision she and I ordinarily attribute to ourselves.

But now what does adoption of the Rosenthalian higher-order thought theory tell us about this scenario? Unless we can find some reason to suppose that the visual judgments we imagined Linda making would lack "visual quality" in Rosenthal's sense, we should have to conclude that she would have everything

she could possibly need for phenomenal vision. She's got the visual quality, and she's got the higher-order thought: thus things would *have* to look somehow to her. Glenda on the other hand, would seem to be missing an essential element. For she would be missing the crucial "higher-order thoughts" that are supposedly needed to go from merely having the sort of visual qualities found in blindsight to things looking some way to a person. To be sure, she (sometimes) has higher-order thoughts of a kind: she thinks, e.g., that it "looks brighter to me *there*," and she thinks that in this sense she sees something there. But those are not the thoughts she *needs*, according to the theory: she needs to attribute to herself a "visual quality," of the sort present when "looking" or visual appearance as she understands this, is absent. So adoption of this higher-order thought theory results in a *curious reversal* of our original scenario. Linda, the blindsighter, whom we were conceiving of as bereft of visual appearance, *has everything she could possibly want for having it*, and Glenda, to whom we said things *do* look some way, couldn't enjoy visual appearances—*she* must be the blindsighter.

But all this is just to say that adoption of the Rosenthal account would render our original scenario inconceivable to us, incoherent: we would have to conclude that each subject both was and was not a blindsighter. We might add: if this theory is supposed to tell us what the "nature" of phenomenal vision is—what "constitutes" seeing phenomenally, it would seem to tell us that either Linda must *have*, or that Glenda must *lack* phenomenal vision—or both—where 'must' has more than merely nomological force—the force, perhaps, of metaphysical necessity. And so, by my test, adoption of this higher-order representationalist theory of consciousness makes visual appearance intellectually disappear.

If these difficulties are acknowledged, it seems that the following three options emerge for the ambitious physicalist. Option one: stick with the theory, and firmly deny that we enjoy visual appearances in the sense that would leave open the intelligibility and, in a broad sense, the possibility, of a contrast between Linda's blindsight and Glenda's blursight. An examination of the reasons one might offer in support of such an attitude is more than I can do here (though I do consider this elsewhere in detail).⁹ For current purposes, I just want to note that a strong justification for this line is needed. If you are going to maintain that scenarios whose conceivability seems to embody a cogent understanding of visual appearance should be regarded (when we adopt the true theory of vision) as downright *logically incoherent or unintelligible*, or that they would then run counter not just to real world psychology, but *transgress the very limits of metaphysical possibility*, then you need to offer some very compelling reasons. And with respect to the latter claim, invocation of stock examples of "a posteriori necessities" plausibly established by thought experiment is not enough to justify the needed analogy.

A second sort of response would be simply to disavow any claim to discover more than nomological necessities here. For example, you do not say that *all we have in mind* when we contrast things looking some way to Glenda,

though not to Linda, is that Glenda can make spontaneous discriminations of a greater range than Linda. Rather more modestly you say that, in the world as it actually is, those to whom things look a certain way can make more and better spontaneous visual discriminations than blindsighters. Or else, you could limit yourself to saying that it is nomologically impossible for blindsight vision to give one's beliefs the "poise" that visual appearance does, or to supply the range of discriminations afforded by any non-conceptual appearances of what is in one's visual field. Or you could confine yourself to saying that, in the actual world, it just so happens that subjects to whom colors or shapes look somehow are also always spontaneously *thinking* at the time that something looks somehow colored or shaped to them.

These more modest claims (especially the last) would not be beyond challenge. But the point I wish to make here is that such modesty would abandon the explanatory ambitions of the original views. For example, we were originally told, in effect, that it would explain phenomenal vision, and reveal its nature to us, if we accepted that one's conscious visual states are continually objects of one's own thoughts. But if one maintains no more than that, wherever phenomenal vision occurs, it is targeted by higher-order thought, then we will lack reason to think that the presence of higher-order thought *explains* the fact that things look somehow shaped and colored to people. We may say: "If anything explains anything here, it's the fact that things *look* thus and so to people that explains the fact that they *think* they do." Similar remarks hold in the other cases. It will be the fact that varying shapes and colors look different to us that explains our talent for discriminating them in judgment—not the other way around. And it will be the manner in which things visually appear that explains why such appearance is non-conceptual in nature (if it is). And if visual appearance nomologically coincides with some functional role, we should conclude that the appearance *plays* the role, not that it *is* the role. To rest content with nomological claims would be to give up the original explanatory ambitions of the theories I've mentioned.

A third response to the challenge I've raised would hold onto the explanatory ambitions, and continue to search out some deeper necessity that will sustain them, but *alter or amend the theory*. For this, one might look to some difference between actual normal human vision and blindsight, on which our contrast between Linda and Glenda seems to be silent—i.e., it does not suggest they would be *alike* in this respect. After all there will be many such differences, since normal human visual subjects and blindsighters differ enormously simply in virtue of the fact that the latter have, to put it crudely, *big holes in their brains*. And in conceiving of Linda and Glenda, no suggestion was made that our subjects did not differ considerably in terms of what was literally in their heads. Now then: perhaps we should point to some such difference to get our story about what constitutes phenomenal vision. We can say of some such internal hidden feature: anyone such as Glenda enjoying blurry limited visual appearances *must have* that, and anyone such as Linda (to whom things look no

way at all) *must lack* it. For that feature, together with other features we can characterize in appearance-free terms just *constitutes* visual appearance. Thus one might hope to sail through my test for phenomenal realism.

If one took this approach, we may ask how far one should go in making the theory's conditions for phenomenal vision dependent on the specifics of our psychology. Do we continue to honor the traditional functionalist idea that mental states are realizable in a very *physically diverse* range of beings? How similar must one be to normal human subjects to have what's minimally sufficient for phenomenal vision? One way to go here is illustrated by Peter Carruthers. His general view of phenomenal consciousness (Carruthers 2000) bears similarities to the two forms of reductive representationalism just considered. For he takes phenomenality to require both a "fine-grained" content, and a layer of higher-order mental representations. But he responds to the challenge posed by hypothetical Linda-style blindsight with an appeal to the idea that in *actual* blindsight, the subject's discriminations (e.g., of shape or orientation) are generated partly by reliance on certain motor pathways in the brain, whose activity serves as a kind of "unconscious cue" to the character of the visual stimulus, disposing the subject to give certain answers to questions about what shape or orientation was found in the "blind" field (Carruthers 2001).¹⁰ This suggestion apparently is that it forms part of the nature of *phenomenal* vision that no such reliance on motor cues is involved. What would constitute the difference between the situation in which the stimulus doesn't look anyhow to the subject, but is discriminated only via spontaneous low-grade blindsight, and the case where it does look somehow to her, if only blurrily—is that in the latter, phenomenal case, discrimination is carried out in the absence of such motor cues.

But even granting that our discriminatory response in phenomenal vision does not actually rely on these motor cues, why suppose this to be anything more than a fact about how phenomenal vision happens to work in human beings? Why suppose there is some deeper, "constitutive" relationship here, of a kind that would tell us it is metaphysically impossible for stimuli to look any way to a subject whose motor pathways contribute in a certain way to her capacity to make spontaneous visual discriminations? The point is not to suggest that there need be *no further* respect—beyond the presence or absence of visual appearance, together with whatever essentially appearance-dependent attitudes one admits—in which two subjects *must* differ, which is necessarily sufficient to make them thus differ phenomenally. The question is—what justifies us in choosing some specific difference as a *constitutive* difference? Maybe, the absence of these motor cues is not constitutively necessary—but at most necessary only in some weaker sense: perhaps it is nomologically necessary for phenomenal human vision. And why could there not be other creatures that have non-phenomenal vision (blindsight) without the use of these motor cues?¹¹ Perhaps then the constitutive difference is to be sought in some *other* aspect of the actual neuropsychological differences between the brains of normal human

subjects and blindsighters. Just how do we distinguish the merely nomologically necessary from the constitutive?

One might say that awarding *some* differences constitutive status would wrongly delimit the range of physical diversity in the possible realizations of phenomenal vision. But once we have admitted that our understanding of the difference between the presence and absence of phenomenal vision is not exhausted by a difference in manifest functional roles, it becomes unclear just what entitles us to draw the bounds of possible realizations in a definite manner. If all we really had in mind when contrasting visual appearance with its absence were some difference in manifest “talents,” then it would be a form of human *neuro-chauvinism* to suppose that some hidden physical aspect of how those talents are realized in us constitutes this difference quite generally. But once this sort of functionalism is exposed as a lack of realism about phenomenal vision, a less liberal approach can no longer be dismissed as chauvinist. At least *some* degree of *conservatism* is called for, if we are to remain realists, while sticking with ambitious physicalism. The problem though is—just how are we non-arbitrarily to select out some particular actual hidden difference and privilege it with “constitutive” or “metaphysically necessary” status?

This brings us finally to the dilemma that I announced at the outset—one that seems to arise for efforts to justify an ambitious form of physicalism. Ambitious physicalists are not content with the supervenience thesis that there can be no mental difference without a difference of a kind that can be captured in the vocabulary of physics. And it would not be enough for them to say that each experiential type (or consciousness generally) is identical with *some physical property or other*. They want to make more specific explanatory claims about the “physical nature” of (for example) vision, committed to more than just nomological necessities. But then they face a dilemma. Either they fail the test for phenomenal realism when they try to absorb visual appearance entirely into something *manifest* and described in *wholly other terms*, or else they search for specific *hidden* features to which they can award deeply necessary or constitutive status. And in the latter case they risk being either too liberal or too conservative in the diversity allowed subjects of phenomenal vision, without justifying a clear standard by which to avoid excess in either direction, so as to ground their claims to the deep necessities they purport to discover.¹²

5. Life without physicalist ambitions

Naturally, more should be said about whether some form of “ambitious physicalism” can resolve the dilemma posed. But if, as I think, the dilemma is deeply rooted, one should ultimately give up on such doctrines. Is the prospect of life without such physicalist ambitions so worrisome? You may be reluctant to abandon them, if you think they offer the only way to honor the success of the natural sciences, or to subordinate philosophical thought properly to empirical research, or to avoid some collection of supposed horrors associated with

dualism. But a robust phenomenal realism need bear no hostility to experimental science just because it wants nothing to do with conceptual and ontological theses vulnerable to the dilemma posed above. Nor is it even clear we need then embrace dualism in any definite sense. At least the argument so far leaves one free to endorse some physicalist supervenience thesis and a “some-physical-property-or-other” identity claim, should that provide any metaphysical comfort.¹³

But maybe it seems that limiting oneself to a modest physicalism, or forsaking physicalism altogether, ultimately means giving up on any scientific explanation of facts about phenomenal experience. This seems to me an exaggerated worry. Without any help from the semantic or metaphysical theses of reductionist philosophy, we may hope to discover systematic enough relationships between what can be judged to occur by first-person reflection, and what can be discovered by probing our hidden features through brain science, that we can answer questions about why, for example, a person is subject to certain chronic feelings of pain, well enough to satisfy what ordinarily motivates such questions. Imagine that we can respond to such questions well enough to have a systematic means of predicting, altering and eliminating in a highly targeted way various sorts of pain in human beings, with reference to fine-grained details in the activity and structure of the central nervous system. Suppose someone then complains that we still have no *real explanation* of why a given subject is or is not feeling pain, because we still cannot say why our world should not have been a pain-free, non-phenomenal “zombie world,” type identical at the level of fundamental physics: “Why is there consciousness rather than mere matter?” But perhaps we will—perhaps we should—find our lack of an answer to *this*, and other similar, residual mysteries, no more deeply threatens the explanatory powers of science than a persistence of the question, “Why is there something rather than nothing?”

Finally, we need not fear that abandoning physicalist ambitions would deprive philosophy of mind of any positive content or aims. For our attempts to discover something systematic in the relationship between what is hidden inside our bodies and what is manifest in our experience are only as good as our understanding of the terms by which we conceptualize the latter. But the quality of that understanding is constantly threatened by obscurities and confusions that can be exposed and dispelled only by the sort of searching, independent-minded critical dialogue that will look for all the world like *philosophy*. And the clarity that may result from such dialogue (however incomplete and tenuous) can be of value not only in framing a conception of what a science of consciousness is out to explain. It can also help us to respond to philosophical questions about knowledge, meaning, ethics, and aesthetics—and, generally, to examine our lives. So there can be more to philosophy of mind than arguing over physicalism.

NOTES

¹ I do not mean to deny any logical distinction among sentences of these forms. For example, ‘It looks to me as if that’s a green one and that’s a yellow one’ does not entail ‘That one looks green to me, and that one looks yellow to me.’ For the first statement may be true in a situation where, in a very dimly lit room, one item merely looks darker to me than the other—and on that basis (plus some background assumptions) I’m inclined to judge the one is green, the other yellow. In that case, it would be wrong to say simply, ‘one looks green to me, the other yellow.’ (I adapt this example from Maund (2003).) However, one can grant such observations, while holding that the use of ‘looks’ in each of the sentences above counts as ‘*phenomenal*.’ It seems I would disagree with Maund, in that I don’t think reports of how something looks to us in a sensory and phenomenal sense have to be confined to a specific grammatical form, and I would not contrast ‘phenomenal’ with ‘intentional.’ Many phenomenal visual features we ordinarily enjoy are in my view inherently intentional.

² Here I take no position on the question of whether only *sensory* appearances (e.g., what’s reported by ‘looks’ in a sensory sense) are phenomenal. Elsewhere (Siewert 1998, Chapter 8) I argue that non-sensory conceptual thought is phenomenally conscious.

³ My discussion here adapts, revises and condenses core aspects of my earlier discussion of blindsight and consciousness (Siewert 1998, Chapters 3-6)—with a focus here on ‘looks’ and ‘visual appearance’ rather than ‘phenomenally conscious.’ What I say (now and before) bears similarities to Block’s (2002) use of blindsight to explain the phenomenal concept of consciousness, but there are notable differences. For instance, I do not rely on his concepts of “access” and “monitoring” consciousness, and my Linda character is not described in the same way as his “superblindsighter.”

⁴ I am deeply indebted to Weiskrantz’s research on blindsight. But it seems to me he goes awry when he talks about consciousness, by suggesting that all his subjects are missing for visual consciousness is some monitoring function—“comment-ability”—with respect to a type of vision that may take either conscious or unconscious forms. Weiskrantz’s discussion encourages one to think of blindsight as a condition in which one sees, but is “blind” to the fact that one does. By contrast, I conceive of it as a condition in which the subject is, in a sense, *blind*, and (in another) *sees*. The subject is “blind”: the stimulus doesn’t look anyhow to him. But he “sees”: he discriminates it via his visual system. What is crucial is to acknowledge the intelligibility of the latter (my) interpretation.

⁵ My talk of “appearance-dependent ways of understanding what is judged” is intended to leave open just how we should interpret, or even whether we should admit, the notion of distinctive “phenomenal concepts.”

⁶ This hypothetical scenario appears to have been partially realized in Weiskrantz’s subject D.B (Weiskrantz 1986). At one stage D.B. reported regaining limited blurry sight in part of his formerly “blind” field. But Weiskrantz found that his discriminations of stimuli in this area were by certain measures actually less acute than his blindsight discriminations.

⁷ Or so I argue in Siewert (1998, pp. 130-3)

⁸ It should be clear that I believe this problem arises whether we take functionalism as a conceptual thesis (as in Armstrong (1968) and Lewis (1972)) or as a “scientific essentialist” thesis about properties (as in Putnam (1967)). Notice also that my argument here is *not* the “absent qualia” argument, if that is supposed to rest on the claim that two subjects who are *completely manifestly functionally equivalent* might nonetheless differ experientially.

⁹ In Siewert (1998, Chapter 5, and 2003)

¹⁰ I respond to Carruthers in Siewert (2001).

¹¹ Carruthers would perhaps reply that some creatures (young children and non-human animals) *do in fact* have non-phenomenal vision without reliance on motor cues, because they lack the disposition (which he thinks essential for phenomenal perception) to think higher-order thoughts about their sensory condition. But this would just makes matters worse, for two reasons. First, it seems he will need to say that what is essential are dispositions to think higher-order thoughts attributing lower order potentially non-phenomenal sensory states (like Rosenthal’s quality states). But things look somehow to *me* even though I am not generally disposed to attribute to myself such blindsight-compatible states. Second, even if (as seems plausible) young children and non-human animals lack cognitive skills necessary for thinking about their own minds, *it does not follow* that they never *feel pain*, and that nothing *looks* any way to them. But, by some kind of necessity, it *would* follow, if Carruthers were right.

¹² One might say that this is another way of raising the general problem originally articulated in Block (1978): how to justify a theory’s claim to be neither excessively “liberal” nor “chauvinistic.”

¹³ However, I am not convinced even modest physicalism can fend off the combined force of the conceivability and knowledge arguments defended by Chalmers (1996). But it also will not be clear to me that anti-physicalism should be a *dualism*, until it is clear to me that the types by which I understand my

actions, and by which I identify perceptible things, can be clearly parceled out into purely mental and purely physical components.

REFERENCES

- Armstrong, D. (1968), *A Materialist Theory of Mind*, (London: Routledge & Kegan Paul).
- Block, N. (1978), 'Troubles with Functionalism', reprinted with revision and abridgement in Block (2007).
- (2002), 'Concepts of Consciousness', reprinted in Block (2007)
- (2007), *Consciousness, Function, and Representation: Collected Papers, Volume 1*. (Cambridge, MA: MIT Press).
- Carruthers, P. 2000. *Phenomenal Consciousness*. (Cambridge: Cambridge University Press).
- (2001) 'Who is Blind to Blindsight?', *Psyche*, 7(04)
- Chalmers, D. (1996), *The Conscious Mind*, (Oxford: Oxford University Press).
- (ed.) (2002), *Philosophy of Mind: Classical and Contemporary Readings* (Oxford: Oxford University Press).
- Dennett, D.C. (1991), *Consciousness Explained*. Little, Brown.
- (1993), 'Living on the Edge', *Inquiry* 36: 135-9.
- Lewis, D. (1972), 'Psychophysical and Theoretical Identifications', reprinted in Chalmers (2002).
- Maund, B. (2003), *Perception*, (Montreal: McGill-Queen's University Press).
- Putnam (1967) 'The Nature of Mental States', reprinted in Chalmers (2002).
- Rosenthal, D. (2002), "Explaining Consciousness", in Chalmers (2002).
- Siewert, C. (1998), *The Significance of Consciousness*, (Princeton: Princeton University Press).
- (2001) 'Spontaneous Blindsight and Immediate Availability: Reply to Carruthers', *Psyche*, 7 (07).
- (2003), 'Eliminativism, First-Person Knowledge, and Phenomenal Intentionality: Reply to Levine' *Psyche* 9(03).
- Tye, M. (1995), *Ten Problems of Consciousness*, (Cambridge, MA: MIT Press).
- Tye, M. (2002), *Consciousness, Color and Content*, (Cambridge, MA: MIT Press).
- Weiskrantz, L. (1986), *Blindsight: A Case Study and Implications* (Oxford: Oxford University Press).
- (1997), *Consciousness Lost and Found*, (Oxford: Oxford University Press).