




PROJECT MUSE®

Does Experience Have Phenomenal Properties?

Geoffrey Lee

Philosophical Topics, Volume 44, Number 2, Fall 2016, pp. 201-230
(Article)

Published by University of Arkansas Press



➔ For additional information about this article

<https://muse.jhu.edu/article/650893>

Does Experience Have Phenomenal Properties?

Geoffrey Lee
University of California, Berkeley

ABSTRACT. What assumptions are built into the claim that experience has “phenomenal properties,” and could these assumptions turn out to be false? I consider the issue specifically for the *similarity relations* between experiences: for example, experiences of different shades of red are more similar to each other than an experience of red and an experience of green. It is commonly thought that we have a special kind of *epistemic access* to experience that is more secure than our access to the external environment. In the first part of the paper, I argue that one way of elucidating this claim is especially plausible—that systematic error, of the kind subjects make about the external environment in traditional “skeptical” scenarios, is not conceivable for introspection of experience, including for our knowledge of similarity relations. I argue that focusing on similarity relations gives us a more interesting version of the argument than for other forms of experiential introspection. Then in the second part of the paper I describe an example, inspired by a similar case due to Sydney Shoemaker, in which a subject, despite being fully rational and attentive, apparently *is* systematically mistaken about the character of their experience in a surprising way. I argue that the example calls into question whether there are properties of experience satisfying the epistemic access constraint, and therefore whether experience has “phenomenal properties” in the intuitive sense.

I. INTRODUCTION

The starting point for much discussion in philosophy of mind involves postulating the existence of “phenomenal properties” or “qualia”—properties of experiences that determine the “subjective character” of the experiences, and whose mysterious nature we would like to understand. But what kind of postulation is this? Could it turn out to be mistaken? What assumptions—implicit or explicit—are involved in the idea of a phenomenal property, and how should these assumptions guide theory building in this area? This paper is an attempt to shed some light on these issues (note: in what follows, I will use “qualia” and “phenomenal properties” interchangeably).

My focus in particular will be on the notion of *phenomenal similarity*. “Phenomenal properties,” if there are such, are commonly thought to form quality spaces organized by the similarity relations between constituent properties (Goodman 1966; Clark 1992); for example, color experiences can be related along at least three dimensions of similarity: hue, saturation, and brightness. The goal of the paper is to argue for two main points about qualia that relate to phenomenal similarity. First, a fairly familiar point: I’ll argue (section 3) that a deeply held assumption about qualia is that we have a kind of epistemic access to them (including their similarity relations) that is in an important respect (elucidated below) *less fallible* than our perceptual access to the environment. As I will explain, I think this is one of the strongest reasons for saying that there are substantial theoretical presuppositions built into the notion of “qualia,” in particular that there are a priori epistemic/functional constraints on our *access* to qualia.¹

On the other hand, and this is the second main point of the paper, I will argue (sections 4 and 5) that in fact experience does not have properties satisfying such functional/epistemic constraints. My argument elaborates a problem Shoemaker (1982, 2006) raised for his functionalist view of phenomenal similarity, later dubbed “Shoemaker’s paradox” by Stalnaker (1999). It involves a hypothetical being who is, arguably, best interpreted as being systematically mistaken about the character of her experience.² Moreover, my argument for the *first* main point of the paper is *also* inspired by an argument Shoemaker (1975, 1981a) gave in favor of his functionalist view. So this paper is very much a tribute to his wonderful writing on this topic, inspired by a belief that there are ways in which his arguments can be significantly improved.

The second main point of the paper could be said to call into question whether experience has “qualia.” I won’t exactly endorse this Eliminativist line, however; rather, I want to conclude that we face an interesting dilemma. Either we should be eliminativists about qualia, or we should accept that they only exist in a somewhat

1. See Kind (2001) for a representative example of a view like this.

2. For some different considerations in favor of the fallibility of introspection, see Schwitzgebel (2008, 2011).

revisionary sense, on which our epistemic access to them is only as safe as our perceptual access to external objects. On this latter horn, it is conceivable that we suffer systematic illusions about the qualitative character of experience; it might even be conceivable that our experience doesn't have qualitative character at all! So my conclusion is: either qualia don't exist, or our access to them is fallible enough that it is conceivable they don't exist!

Before we get into the main business of the paper, I want to first make a few general points about what it means to debate the existence of qualia (impatient readers can skip to section 3).

II. QUINING THE QUALITATIVE

Is it possible to reasonably doubt that qualia exist? On the one hand, this can seem crazy, because it can seem tantamount to denying the existence of conscious experiences. And what could be more certain than the existence of conscious experiences? On the other hand, there is a venerable tradition, going back at least to Dennett's "Quining Qualia,"³ of saying that there is something dubious about the notion of "phenomenal property" or "qualia," and that perhaps we should stop believing in them.

We could define "qualia" to be whatever features individuate an experience, such as the properties of an experience that make it a color experience rather than a shape experience, or which distinguish one kind of color experience from another. Thought of this way, qualia exist if and only if experiences exist. We could proceed by defining "qualia" this way⁴ (although see note 4 for some concerns about this), but I agree with Dennett (1991) that we typically have something a bit more specific and a bit more loaded in mind.

Dennett is often accused of attacking a trumped-up notion of "qualia" that no one really ever believed was applicable. These days, philosophers typically believe that they can operate with a less theoretically loaded notion—what Keith Frankish (2012) has dubbed the "diet" notion of qualia, to contrast it with Dennett's target, "classic qualia."⁵ But I don't think it's obvious that "qualia," even when used in a deliberately "diet" way, is completely free of substantial presuppositions, as I will

3. Rey (1983) is a notable precursor to Dennett's influential (1988) discussion.

4. One problem with this approach is that some individuating features of an experience might be aspects of neuro-functional states that are necessary but not sufficient for an experience to exist, and which are not themselves introspectively accessible to the subject. We could get around this by stipulating that we are talking about maximal individuating features that are sufficient for the existence of an experience. But the problem with this is that we typically want to count as phenomenal properties nonmaximal aspects of experience that might not be able to exist without other phenomenal features. How do we individuate these except in terms of what is epistemically accessible?

5. A related distinction is Levine's (2001, ch. 5) distinction between "bold" and "modest" "qualiophiles."

argue. Indeed, it's not clear whether the notion can be completely rid of such presuppositions without being rendered completely empty.

This latter worry suggests a kind of “conceptual eliminativism” that might be pressed by a qualia-skeptic: that there is no distinctive “diet” notion of qualia (this is Frankish's [2012] position).⁶ But here I'll be presenting a more Dennettian line: the notion itself may be cogent, but it contains presuppositions that may not in fact be met. The difference with Dennett is that the presuppositions I'll consider, to do with our epistemic access to qualia, are of a more moderate and plausible kind than those he went after.

It is worth mentioning up front one reason for qualia-skepticism that won't be relevant here. Some argue against the existence of qualia on the grounds that experience is “transparent”—in particular, on the grounds that attempts to introspect qualia only reveal (allegedly) mind-independent properties like colors, tastes, smells, and other perceptible features; introspection does not reveal “intrinsic” features of experience (Tye 2002; Harman 1990). This argument doesn't move me; experiences could be constituted by experiential relations between subjects and mind-independent properties; in which case the “intrinsic” properties of the experience would involve the ways in which it relates the subject to these mind-independent features, and awareness of external features would be a way of coming to know about the intrinsic character of the experience (Siewert [2004] and Levine [2007] both make similar responses to transparency arguments). In short, we can use “qualia” in a sufficiently “diet” way that their existence is compatible with such a “relational” view.

What the point does bring out, however, is the way in which the notion of “qualia” is connected with the theme of this volume: qualia (at least in stereotypical perceptual cases), involve the way external objects and events *appear* to us in experience, and we can come to know about them by attending to these external appearances, which may (or may not) involve experiencing objects as having mind-independent features. Nothing in what follows will turn on whether such a “relational” understanding of experience is correct, however; in particular, what I say is *also* compatible with experiences *not* being constituted by a mental relation to external features; on such a view, the phenomenal appearances of objects might

6. Frankish (2012) contrasts “classic qualia,” “diet qualia,” and “zero qualia.” “Zero qualia” is something close to a purely functional concept: they are “the properties of experiences that dispose us to judge that experiences have introspectable qualitative properties that are intrinsic, ineffable, and subjective.” Frankish thinks that attempts to elucidate the notion of “diet qualia” either enrich so much that it becomes a version of the “classic qualia” concept, or deflate it to something purely functional, similar to “zero qualia.” His argument is that proponents of a concept of diet qualia can only conceive of diet qualia as properties of experience that are targets of phenomenal concepts; but that amounts to picking them out by a functional description like “property causing the use of a phenomenal concept.” The argument rests on use-mention confusion: subjects can also simply *use* phenomenal concepts to pick out qualia, and these concepts are not functional concepts. Therefore, subjects do have nonfunctional concepts of diet qualia (although for reasons discussed in this paper, these concepts might not be entirely theoretically neutral).

be constituted by their disposition to produce (nonrelational) phenomenal states in us, rather than apparent mind-independent features.

To clarify what's at stake here further, it's helpful to introduce some terminology. Call an assumption about phenomenal properties **substantial**, if it is an assumption that can play a theoretical role in guiding our inquiry into the nature of experience—it's the kind of thing that can help us get traction linking experiential facts with neuro-functional facts, for example. I'll assume that substantial assumptions are those that have functional implications: they constrain the functional structure of the subject who has the phenomenal properties. We can use this notion to contrast **empiricist** and **rationalist** views of our phenomenal concepts. Empiricists think that there are no substantial assumptions built into phenomenal notions; rationalists think that there are. Views will also come on a spectrum from more empiricist to more rationalist.

At the far end of the spectrum of rationalist views is analytical functionalism, the view on which experiential properties are functional properties, and their functional definition can be known just by grasping concepts of them. Since this view has well-known problems,⁷ a more interesting form of rationalism to consider is one on which there are significant a priori constraints that are associated with phenomenal concepts, but it is not the case that these concepts have a full analysis in functional terms. I think this is an attractive view: in part, because it is unlikely to be susceptible to standard problems for analytic functionalism, and in part, because it has independent motivations and benefits (some of which I will describe below). We might call this *Moderate Rationalism*. A specific version of this view, on which there are a priori functional/epistemic constraints on phenomenal similarity, will be my subject here. I'll call this *Moderate Similarity Functionalism*, or "Similarity Functionalism" for short.

For a rationalist, it's also helpful to distinguish *fallible* and *infallible* assumptions. Suppose we decide that condition X on experience is built into the notion of a phenomenal property. Could it turn out that nothing meets condition X, and so nothing has phenomenal properties? If that is conceivable, then we can say that X is a *fallible* assumption. Faced with evidence that X is not satisfied, we might revise our concept of phenomenal property, replacing it with something weaker, or something with different built-in assumptions. If X being unsatisfied is simply inconceivable however, we'll say that it is an *infallible* assumption. "Phenomenal Property" might have both fallible and infallible assumptions built into it—in which case we can imagine being forced to weaken the concept in some ways, but not others. A *fallibilist* is a rationalist who thinks that fallible assumptions are built

7. In particular, the view implies that once you know the functional structure of a brain, you can infer a priori exactly what experiential states are realized in it. However, like many philosophers, I do not believe that *anything* about consciousness can be inferred a priori from knowledge of functional structure, because "functional zombies"—beings with our functional structure but who lack consciousness—are conceivable.

in to “phenomenal property.” Likewise, we can also talk about the *infallibilist* view (which is not incompatible with fallibilism). A *skeptical fallibilist* is a fallibilist who believes that the fallible assumptions are not met (I already mentioned Dennett’s [1988] discussion in “Quining Qualia,” which is fairly clearly an instance of skeptical fallibilism. One other group of philosophers who can reasonably be interpreted as skeptical fallibilists are disjunctivists [see note 8]).⁸

A further important distinction is between *absolute* and *idealized* functional constraints. An absolute constraint is one that applies in any situation, regardless of background conditions. An idealized one requires that certain background conditions obtain. Most salient here are *epistemic* constraints on phenomenal properties that have functional implications only relative to assumptions about the *rationality* of the subject. For example, perhaps an ideally rational subject would, on having an experience, be disposed to classify it a certain way; but with less than complete rationality, a subject may fail to have this disposition, even though they *ought* to have it. In the next section we’ll look at some considerations in favor of rationalism; it’ll be helpful to think about what kind of constraint they really support.

A final preliminary point: Philosophers who are skeptical about whether there is a sharp analytic/synthetic distinction may think that the rationalist/empiricist distinction here is not very important, because there is no important distinction between revising deeply held beliefs about phenomenal properties that are not a priori, and revising our concept of phenomenology to accommodate the failure of a priori assumptions. I’m not completely unsympathetic with this, but I also think that it is too quick a dismissal. Certainly, what I really care about here are the “deep” assumptions about phenomenal properties that guide our thinking in this area, and for some purposes it may not matter whether “depth” involves a priority. But note that if we are infallibilists about certain assumptions, then this most definitely is a kind of rationalism, not empiricism. Furthermore, even in the case of fallible assumptions (such as the kind of link between indistinguishability and phenomenal similarity rejected by disjunctivists), it might be illuminating to treat these assumptions as conceptual: in particular, stripping them away might leave us without a concept that has any determinate content, pushing us toward a Frankish-style conceptual eliminativism. In fact,

8. Disjunctivists hold that indistinguishable perceptions and hallucinations do not have phenomenal properties in common, because there is no experiential psychological state that is common to hallucination and perception (see Campbell 2002, Martin 2006, Hinton 1967 for classic statements of the disjunctivist view). A common reaction to this view is to find it to border on incoherence—isn’t it part of our basic idea of a hallucination that it feels the same as a perception, and therefore has the same phenomenology? The disjunctivist will reply that we are only entitled to assume that these states ‘feel the same’ in the sense that they are *indistinguishable*—which doesn’t imply that they share any experiential core (Martin 2006). But even if that’s right, it might be fair to say that the view does enough violence to our notion of phenomenology (in the way that phenomenal sameness/difference is so sharply dissociated from distinguishability), that it turns out that there aren’t really phenomenal properties if it is true. And indeed, disjunctivists often say that they are skeptical of all this “phenomenal property” talk.

one of the most significant challenges for strongly empiricist views is to explain how phenomenal concepts can be theoretically neutral without being defectively indeterminate. (This will become relevant later, although I won't have space to discuss it in great detail.)

The game plan: in the next section, I give a positive argument for similarity functionalism, a form of moderate rationalism. Rationalism doesn't entail that experience actually has properties that satisfy the relevant a priori constraints, however; in section 4 I give an argument that casts doubt on this.

III. RATIONALISM AND SKEPTICISM-PROOF ACCESS TO EXPERIENCE

In this section, I make a positive case for the version of moderate rationalism I just mentioned, similarity functionalism.

Start with the idea (already mentioned above) that phenomenal properties stand in relations of phenomenal sameness and similarity, and that these relations organize the properties into families, commonly known as “quality spaces” (Goodman 1966; Clark 1992). For example, color experiences are similar along dimensions of hue, saturation, and brightness, which organize them into color space. On a “structuralist” view, phenomenal properties are fully individuated by their location in a quality space (see Clark 1992; Rosenthal 2010, for examples). But even if we reject this view, and hold that each property has an “intrinsic nature”⁹ that transcends its structural relations to other properties, we might well think that these structural relations are determined by, and so reflect, the nature of phenomenal properties: phenomenal similarity is an *internal* relation between phenomenal properties, which partly individuates them. Furthermore, if there are functional constraints associated with phenomenal ascriptions, it's more plausible that these relate to sameness and similarity relations than with whatever intrinsic phenomenal nature is left over. Notoriously, such structure-transcendent intrinsic natures seem to us, at least considering the matter a priori, to also be *functionally* transcendent; for example, in inverted qualia cases we appear to have functional duplication, duplication of phenomenal sameness/similarity, but inversion of the intrinsic natures of qualia.

The fact that attempts to construct inverted qualia scenarios almost always involve trying to preserve phenomenal similarity structures while inverting intrinsic phenomenal natures tells us something interesting about how we intuitively think of phenomenal similarity: we *do* regard it as placing functional constraints; we're

9. This isn't to believe in phenomenal properties as “intrinsic” in some theoretically contentious sense. It is merely to say that there is more to a phenomenal property than its location in a similarity space, a claim that is consistent with most popular views of phenomenology (for example, it is perfectly consistent with externalist intentionalist views).

not at all sure that we can coherently imagine it varying independent of functional structure.

It's precisely this kind of consideration that led Shoemaker (1975, 1982) to believe that phenomenal similarity was a priori definable in functional terms. Roughly, his proposal is this: experiences are similar in a certain respect just if we are disposed to judge they are similar in that respect. We can think of this as a generalization of the highly tempting idea that phenomenal properties are exactly the same just if experiences that have them are *indiscriminable*.¹⁰ The more general idea is that there is no gap between how the facts about phenomenal sameness and similarity are disposed to strike us, and how they actually are: a very strong kind of "special access" constraint.¹¹ For example, on Shoemaker's view, if you are having an experience of three lights, and one appears exactly intermediate in brightness to the other two, then you are disposed to judge that it is intermediate in brightness.

Even though Shoemaker didn't think that an a priori functional definition of an *individual* phenomenal property could be given (more on this below), his view is nonetheless best regarded as a form of analytical functionalism; he thinks that both whether a being is conscious and the similarity structure of a subject's experience are derivable from functional information (what can't be functionally derived are the specific kinds of experiences that stand in these relations). The kind of moderate rationalist I'm interested in need not even go this far; they could

10. Consider the fact that people often talk of "indiscriminable experiences" when they mean "experiences of the same phenomenal kind." Skeptical fallibilist disjunctivists might say that it is this idea that generates intuitive resistance to disjunctivism, because on their view indiscriminable experiences (hallucinations and perceptions) need not share mental properties. Insofar as disjunctivists can be read as rejecting phenomenal properties, note that this actually supports a Shoemaker-style analysis.

11. An important complication in Shoemaker's functionalist account should be noted. For him, the primary functional impact of phenomenal similarity is on externally directed perceptual judgments, not introspective judgments: for example, the phenomenal similarity between color experiences will cause a corresponding judgment of similarity concerning external colors. However, Shoemaker thinks that we make introspective judgments of phenomenal similarity simply by redeploying such perceptual judgments (an assumption we might reasonably call "transparency"); for example, we judge the similarity between color *experiences* by judging the similarity between color *stimuli*. So there is a kind of equivalence between the functional impact phenomenal similarity has on externally and internally directed judgments. This 'transparency' assumption deserves extended discussion, but it won't be my main focus here, even though some version of this principle might be regarded as a further pre-theoretical connotation of phenomenal ascriptions, and might need to be assumed in order to justify the kinds of functionalist claims about phenomenal similarity I'll discuss here. I'll just assume that it's acceptable, so that we can move back and forth between talk of external and internal discrimination.

This link between internal and external discrimination also suggests a picture on which phenomenal similarity between experiences corresponds to the *representation* of similarity between external stimulus features: phenomenal similarity carries information about external similarity. Seen through this lens, Shoemaker can be seen as analyzing phenomenal similarity in terms of informational access: more specifically, in terms of access to *similarity information* about external stimulus features represented in experience. (Those like Block [1995] who have argued for a strict separation of "phenomenal" and "access" notions of consciousness may see this as problematic. I'll return to this below.)

think that there are functional *constraints* associated with phenomenal similarity without thinking that phenomenal similarity can be straightforwardly *defined* in functional terms. That is, there is some kind of a priori link between phenomenal similarity and the functional/epistemic relations that hold between experiences when we make discriminatory judgments of similarity between experiences or external stimuli. This is Similarity Functionalism.

Our intuitive thinking about inversion cases already supports this view, as mentioned. A further important argument in favor of it (although not one I think we should endorse) is what I'll call the *methodological* argument. It is important to understand that the question of how phenomenal similarity and discriminatory dispositions are linked is not just philosophically important, but also has great relevance to the science of perception. It is really the main foundational question in psychophysics, where it is still not clearly understood how to best infer a psychometric function (in particular, the function linking an experiential quality space to an external stimulus space) from subjects' discriminatory behavior; to do this, we need to know how to make the inference *from discrimination to experience*, which is exactly what is at issue here. Fechner proposed that the similarity of experiences corresponds directly to the discriminability of stimuli: experiences typically caused by stimuli that are equally discriminable, in the sense they are equal numbers of just-noticeable-differences away from each other, are equally different. But Stevens, for good reasons, called this view into question, and ever since a debate has raged about how to infer perceptual psychometric functions (which can often be understood as modeling phenomenal similarity) from different kinds of discrimination behavior (see Gescheider 1997 and Laming 1997 for reviews).

It is hard to see how we could adjudicate this foundational debate—or more generally, how there could be any way of figuring out the neuro-functional basis of phenomenal similarity—without having some pre-theoretical grip on how phenomenal similarity affects discriminatory judgments.¹² A moderate rationalist might use this to argue that there must be a priori functional constraints on phenomenal similarity. However, although it is clear that such constraints would have an important role to play in empirical investigation *if* they exist, I think this methodological argument is just an exercise in wishful thinking. It may be that all we have to go on are precariously a posteriori pre-theoretical assumptions. The functionalist is better off finding an independent motivation for his view, and then applying the results to the interpretation of psychophysical experiments.

That said, I now want to consider what I think is the most powerful way to motivate similarity functionalism. I'll call this the "anti-skeptical argument"; it also helps clarify what kind of constraint we are talking about.

12. This can be seen as a specific instance of the kind of methodological problem for consciousness studies discussed in Block (2007) and Chalmers (1998)—a subject's discriminations can be seen as a form of access-consciousness with respect to information about the similarity of stimuli.

This argument is inspired by an argument that Shoemaker (1975, 1981a) himself gave for his view. Shoemaker was trying to argue for analytical functionalism. His argument (at least, one version of the kind of argument he had in mind [see note 13 for exegetical discussion]),¹³ was that any weaker view leads to the conceivability of an absurd form of phenomenal skepticism; in particular, it leads to the conceivability, for any conscious being, of their *not* being conscious. The reason is that if analytical functionalism fails, then functional zombies are conceivable. And according to Shoemaker, functional zombies are convinced that they are conscious. But if it is possible for an individual to be convinced that they are conscious even if they aren't conscious, how do I know that I'm not myself in that predicament?

For reasons given by others (outlined below), this argument probably doesn't work. However, I think there's a better argument in the same spirit. Shoemaker's guiding idea was that certain skeptical hypotheses about consciousness are absurd, and that their absurdity suggests that it is part of our concept of consciousness that we have some kind of special epistemic access to consciousness that allows us to rule these hypotheses out.

That is, the nature of experience puts limits on how badly wrong we could conceivably be about our current experience (such as about whether it exists), and, therefore, it constrains the kinds of (rational) judgments we might be disposed to make in response to the experience: a kind of functional constraint.

There is a plausible version of this style of argument for the phenomenal sameness and similarity relations between parts of your experience. Consider the current experience you're having. If it's anything like mine, it involves a fairly richly variegated array of experience-types: experiences of different shapes, textures, color, sounds, etc. Could your beliefs about this experience be completely wrong? In particular, could you be completely wrong about the kinds of phenomenal

13. Shoemaker's official argument invoked the causal theory of knowledge to argue that if functionalism isn't true, we don't know we are conscious because we aren't causally sensitive to our phenomenal features in the right way; a zombie-state could play the same causal role as a phenomenal state, including producing the same beliefs, and so phenomenal features don't have the right kind of causal relevance in producing *our* introspective beliefs. To this it can be replied (see Block 1980) that the mere possibility of such zombie states doesn't make phenomenal features causally irrelevant to belief, so, on the causal theory of knowledge, we can know about our phenomenal features. But it's clear from Shoemaker's response (1981), that he sees this as missing the point. The driving idea is that it's absurd to suppose that you can't tell the difference between being conscious and being a zombie. A charitable way of construing this is as saying that it's absurd to suppose that your epistemic situation doesn't rule out phenomenal skepticism. Shoemaker, perhaps mistakenly, developed this thought in terms of causal sensitivity of introspective belief, but there may be other ways of understanding your "epistemic situation": for example, in terms of what is conceivable for you, or what you have definitive justification for ruling out. Importantly, my beliefs could be caused by phenomenal states even if (perhaps absurdly) I can't rule out phenomenal skepticism. Hill (1991, ch. 6) interprets Shoemaker's point in this broader way, and Chalmers (1996, ch. 5) and Levine (2001, ch. 5, 6) discuss this broader epistemological puzzle.

similarity relations that obtain between parts of your experience? For example, could it be that you are actually experiencing all the surfaces in your environment as slightly different shades of orange, even though you believe they have very different colors? The proponent of the argument says that such a *phenomenal skeptical scenario* is simply inconceivable. Call this *anti-skeptical rationalism*.

But if anti-skeptical rationalism is true, it is a priori that there are limits to how far wrong certain kinds of introspective discriminations can go. And that can only be the case (so the argument goes) if there are limits to the kinds of (rational) introspective discriminations you are *disposed to make* when you have an experience of a certain type—that is, it can only be the case if there are (a priori) functional constraints on phenomenal similarity. So anti-skeptical rationalism supports similarity functionalism. This, in brief, is the anti-skeptical argument.

Note that here the skeptical absurdity is not generated by the possibility of a functional zombie, as in Shoemaker's argument, but rather the possibility of a *conscious* being who is drastically wrong about the *kind* of experience they are having at a particular moment of time. Also the proponent of the argument is not arguing that we have infallible or super-reliable access to our experiences. It is consistent with the argument that in various ways our judgments about our experiences can go wrong—even quite badly wrong. But the proponent does think there are limits to this—in particular, that there is a kind of fallibility that *perceptual* judgments have that introspective judgments do not have. In the case of perceptual judgments, it is perfectly conceivable in a given instance that these judgments are completely wrong. For example, it might appear to me as if I am surrounded by a coffee shop, when in fact I am a brain floating alone in an infinite void. The idea here is that analogous forms of experiential skepticism are not conceivable in the same way, and in that sense we have “special access” to experience.

The standard response (Chalmers 1996, ch. 5; Hill 1991, ch. 6; Levine 2001, ch. 5, 6) to Shoemaker's anti-skeptical argument on behalf of those who reject analytical functionalism is to claim that our epistemic situation is not on a par with that of the zombie. The fact that we are conscious and they aren't puts us in a superior epistemic situation (although see Balog 1999 and Lee 2013 for some pushback against this). A typical way to develop this response is in terms of two ideas: (1) zombies don't have the concept of consciousness, or concepts of particular kinds of consciousness, so they can't believe that they are conscious; (2) conscious states are immediately available to us for us to think about using phenomenal concepts, and the way we access them with these concepts means we can't be completely wrong about their existence; furthermore, the correct story here doesn't require that analytic functionalism is true. The idea is that (1) distinguishes us from zombies and (2) explains our peculiar epistemic connection with our conscious states in a way that avoids functionalism.

A version of this response might work against Shoemaker, but it doesn't easily generalize to the version of the argument I outlined. In particular, defenses of (2) tend to guarantee only that *individual* ascriptions of phenomenal properties

cannot go wrong—they do not speak to the issue of phenomenal *comparisons*. For example, a fairly plausible story about why individual ascriptions can't go wrong is that uses of phenomenal concepts literally involve the experiences whose properties are being referred to—so that the concepts could not be used unless in the presence of an instantiation of the property being ascribed (Chalmers 2003). Clearly such a story does not require appeal to any form of functionalism—it's not that the experience is guaranteed to *cause* the relevant introspective judgment; it's that the judgment simply could not exist without the experience. By contrast, a guarantee that separate uses of two *different* concepts referring to different phenomenal properties will be aligned in a certain way in a similarity judgment is, by its very nature, a *functional* constraint. So it's much more plausible that the anti-skeptical argument supports a moderate form of functionalism about phenomenal similarity than it is that Shoemaker's similar argument supports analytical functionalism; the same anti-functional maneuver doesn't work in this case.

This is the anti-skeptical argument for similarity functionalism. What responses might an opponent give to it?

First, the distinction I made earlier between *absolute* and *idealized* functional constraints is salient. Does the argument entail that systematic introspective error is absolutely impossible, or only impossible relative to an assumption of rationality (or other idealization)? If only the latter, then perhaps anti-skeptical rationalism only entails a fairly weak form of functionalism. We can distinguish two questions here. First, if a person was sufficiently irrational, could they have radically mistaken beliefs about the nature of their experience? Second, if a person does have radically mistaken beliefs about their experience, does that imply that they are irrational? These are important questions: but it will make sense to delay addressing them until we've considered a relevant case where an individual is, or appears to be, radically mistaken (section 4).

Second, if the argument succeeds, it suggests an a priori access constraint on experience: a subject's access to the similarity structure of their experience is skepticism-resistant. But many philosophers are skeptical of whether there are access-constraints on phenomenal consciousness. Ned Block (1995) famously argued that we should make a clear separation between two concepts of consciousness—the concept of phenomenal consciousness and the concept of access consciousness, concepts which he claims are easily confused. It might be objected that the anti-skeptical argument trades on a confusion between these notions. Access consciousness is an explicitly functional concept, so no surprise it implies functional constraints on access. Phenomenal similarity, it might be claimed, doesn't imply *anything* about access to similarity however.

More specifically, you might think of it like this: phenomenal similarity corresponds to the experience of similarity relations between external properties (see note 10). One's discriminatory judgments are the form that "access" to this similarity information takes. A defender of the phenomenal/access distinction might then take the line that phenomenal consciousness of similarity does not imply access

consciousness of similarity (because phenomenal consciousness does not in general imply access consciousness). So phenomenal similarity doesn't imply anything about a subject's discriminatory dispositions, contrary to similarity functionalism.

Obviously the strength of this argument depends on the strength of the considerations put forward by Block for a sharp phenomenal/access distinction. We need to be careful to distinguish two issues here, however: first, whether there is a *conceptual* access/phenomenal distinction; second, whether phenomenal and access consciousness are *in fact* different phenomena, even if they are conceptually distinct. Conceptual distinctness is more relevant to the debate about rationalism and is far less controversial.

The point I want to make about this here is that the anti-skeptical argument problematizes a certain way of thinking about phenomenal/access conceptual distinctness, a way of thinking that underpins the Blockian argument just given. It seems reasonable to give the following response: even if we accept conceptual distinctness, surely it's still plausible that individuals cannot be completely wrong about their current experiences, and so they *do* have a special kind of access to it. You couldn't rationally think that you are experiencing a room full of objects, even though you are really experiencing a featureless ganzfeld. The kinds of considerations that favor conceptually distinguishing phenomenal and access consciousness have no force against this response, it seems to me. So where does that leave us?

I think that what it suggests is that, even if the concept of access consciousness, as it is defined by Block, is separate from the concept of phenomenal consciousness, that doesn't mean that there aren't access-like functional constraints built into the notion of phenomenal consciousness.

There is an ambiguity in the notion of "access consciousness" that is relevant here. On the one hand, it can be understood as a purely functional notion whose home is more in cognitive science than folk psychology. An example of a notion of this kind is Baars (1997) notion of "availability of information to a global workspace." It's obvious that such a technical functional notion is distinct from the notion of phenomenal consciousness, and also that phenomenal consciousness does not a priori entail access in this sense. But, it's also obvious that no one would confuse these notions—and Block introduces access consciousness by claiming it is easily confused with phenomenal consciousness.

On the other hand, the notion of access in play might be a more folk-psychological one, a kind of access that *is* familiar to anyone who has it. It's the kind of access we typically have to our phenomenal states: we might call this "phenomenal access." I think it's not at all clear that concepts of phenomenal states don't entail *constraints on phenomenal access*. The constraints could be weak enough to allow for the kinds of inaccessible experiences that Block has advocated for. But at least, they would require that certain kinds of judgments comparing experiences we *do* have phenomenal access to could not be radically mistaken. Given our strong a priori confidence that phenomenal skepticism is inconceivable,

the moderate rationalist is on fairly firm ground asserting that there *are* such constraints, and I do not see anything in Block's case for a phenomenal/access distinction that cuts against this. In short, even if "access consciousness" (however it is understood), and "phenomenal consciousness" are different notions, that doesn't mean that "phenomenal consciousness" is an *access-free* notion: it may well entail certain things about access, understood as *phenomenal access*.

A third important objection. The argument is supposed to show that certain introspective beliefs are skepticism resistant: but does this tell us something about the nature of *experience*, or rather about the nature of *beliefs about experience*? In particular, maybe the inconceivability of phenomenal skepticism is to be explained in terms of conditions that have to be in place for us to have beliefs *about* our phenomenal states: that is, meta-semantic constraints on phenomenal *concepts*, not functional constraints on phenomenal *states*. Compare how, on many meta-semantic views, we can't refer to, or think about, external objects unless we are causally connected with them in the right way: that doesn't imply that it is in the nature of external objects that they are causally connected to our beliefs. Why not say something similar here?

In response, it's not clear that the similarity functionalist should see this as an objection. They care about a priori functional constraints on phenomenal similarity: whether the constraints reflect the nature of phenomenal properties or phenomenal concepts is not necessarily something they are taking a stand on. For example, they could (but they need not), give their anti-skeptical argument a meta-semantic spin. They might ask: is the successful reference of phenomenal concepts compatible with the truth of phenomenal skepticism? If my current experience is nothing like the way I think it is, because its similarity structure is completely different from how I think it is, can I nonetheless successfully think about individual aspects of the experience? A negative answer might suggest the following argument for anti-skeptical rationalism:

- (1) It is inconceivable that my current uses of phenomenal concepts fail to refer.
- (2) If I am in a phenomenal skeptical scenario, then my current uses of phenomenal concepts fail to refer.
- (3) Therefore, I am not in a phenomenal skeptical scenario.

(2) could be supported in terms of a fairly plausible *contrast constraint* on phenomenal reference: to refer to a phenomenal property, I have to have some knowledge of how it contrasts with other salient phenomenal properties in my experience (much as singling out a physical object for reference requires at least a somewhat accurate sense of how it contrasts with its surroundings: where its physical boundaries are, etc.). For example, suppose you are having an experience as of a red square in between two blue squares. Suppose you believe that the red square has a very similar color to the blue squares (which might make you disposed to judge that it is blue). Can you nonetheless form a phenomenal concept of the red shade that is in your experience? This seems dubious.

We could appeal to some such constraint to support the above argument. On the other hand, the constraint could be used in a different way. We could invert the explanation, *starting* with the idea that phenomenal skepticism is inconceivable, and use this to explain why one possible source of reference failure of phenomenal concepts can be ignored. On this alternative story, the anti-skeptical constraints have meta-semantic *repercussions*, but may not be meta-semantically grounded. To my mind, the second story is more plausible, but, again, the moderate rationalist need not take a stand on that here. (The question of whether phenomenal skepticism threatens our ability to think about phenomenal properties will be relevant again later, however.)

To sum up—I think the anti-skeptical argument is a forceful way of explicating the idea that is part of our conception of phenomenology that it is accessible to us in a certain special sense—in particular, it is accessible in a way that rules out phenomenal skepticism. This is the anti-skeptical rationalist view. And anti-skeptical rationalism is a plausible way of motivating the idea that there are indeed a priori *functional* constraints associated with phenomenal ascriptions: the similarity functionalist view.

However, as noted earlier, it is one thing to say that we conceptualize the properties of experience in a certain way, and quite another to say that the properties of experience actually fit our conception. To avoid confusion here, it will be helpful to distinguish between *conceptual* similarity functionalism and *metaphysical* similarity functionalism. Conceptual similarity functionalism is the conceptual thesis I just argued for; metaphysical similarity functionalism is the view that the properties of experience actually fit our conception. In the next section, I will argue that a phenomenal skeptical scenario is in fact possible after all, suggesting that MS functionalism is false. This raises the question whether experience even has any phenomenal properties, to which I will turn in the last part of the paper.

IV. AGAINST METAPHYSICAL SIMILARITY FUNCTIONALISM

Before getting to the main argument I want to consider here, let me just briefly mention another reason I am skeptical about MS functionalism that deserves a separate extensive discussion. We so far have said very little to make more precise what the alleged functional constraints on phenomenal similarity would look like. Shoemaker assumed that there will be a simple a priori relation between phenomenal similarity and a subject's dispositions to make similarity judgments: if a certain similarity relation obtains, then the subject is disposed to judge that it obtains. But the history of empirical work on psychophysical scaling suggests quite the opposite; it is a highly nontrivial problem to infer a psychometric function (which can often be regarded as modeling a subject's experience) from discriminatory judgments, and the problem is precisely that of determining the relationship

between the properties of experience and dispositions to make discriminatory judgments. There are many different kinds of discriminatory judgments subjects can be asked to make in response to a stimulus of a certain kind, such as different shades of color or sounds of different intensities: there are simple pairwise discriminations, interval bisection tasks, direct magnitude estimation tasks, just to mention a few. And for any one of these there are different interpretative principles we could adopt for inferring the underlying experiential similarity facts. Different discrimination tasks and interpretative principles can deliver totally different verdicts. To return to the example of the Fechner and Stevens debate: they notoriously disagreed on whether psychometric functions typically have a log or exponential form, basing their views on quite different kinds of discriminatory judgments (magnitude estimation vs. discrimination threshold tasks) and interpretative principles. The debate about this continues (see Laming 1997; Lockhead 1992; and Shepard 1978 for more discussion). If it is resolvable at all, it will be on partly empirical grounds, suggesting that at best there are abstract a priori constraints on a solution, not a simple a priori relation between phenomenal similarity and dispositions to similarity judgments.

Important as these problems are, however, here I want to focus on a more purely philosophical argument against MS functionalism, to which I now turn.

I want to present an argument against similarity functionalism that hinges on a dilemma. The functionalist can develop his view either as a version of *functionalist structuralism* or as a version of *role-realizer functionalism*. But the first view is very implausible, and the latter, although prima facie attractive, turns out to be incoherent (or at least incoherent in combination with plausible assumptions), for reasons I will explain in some detail.

Start with functionalist structuralism. As mentioned above, structuralism is the view that phenomenal properties are fully individuated by their similarity relations to other phenomenal properties: they are points in similarity structures. The functionalist version of this view says that these similarity relations have functional definitions (or at least are functionally constrained); presumably these have to do with discriminatory dispositions of the kind mentioned above. Shoemaker (1982) rejects structuralism because he believes in the possibility of inverted spectrum cases, which involve individuals with structurally identical yet qualitatively different spaces of phenomenal properties. But one needn't appeal to such a controversial possibility to see that the view is highly problematic.

Surely such structural relations alone are just *too abstract* to individuate experiences. There are several ways of developing this objection, but here's one illustrative problem: many extensive magnitudes that we can perceive—length, duration, weight, etc., have isomorphic one-dimensional quality spaces—but nonetheless are clearly qualitatively distinct; quality spaces are not individuated by structure alone. (It's also unclear whether a qualia's position in a space with a certain structure is *necessary*: can't quality spaces sometimes be augmented [e.g., through enhancing or adding new perceptual capacities to make a larger space], or diminished, without completely changing the character of the experiences in the space?)

Structuralism nonetheless has its adherents (Clark 1992; Rosenthal 2010) and probably deserves more extensive discussion. But here I want to focus on what is probably the more *prima facie* attractive alternative, some version of *role-realizer functionalism*, on which both the functional role of a phenomenal property and the neurofunctional realizer of the functional role are relevant for determining what kind of phenomenal property it is. More specifically, as on a functionalist structuralist view, the potential for the subject to have experiences with a certain similarity-space structure (such as the space color experiences) will be determined by a structure of functional relations between neural states. But since on this view the nature of the phenomenal properties transcends their similarity relations (they have a transcendent “intrinsic nature”), we need to fix further facts about how this functional structure is physically realized to determine which specific phenomenal characters stand in these similarity relations (for example: is this a normal color perceiver, or a color perceiver with inverted qualia?).¹⁴

I’m going to argue against role-realizer functionalism using an example based on a similar example constructed by Shoemaker (1982). The idea is that the RR functionalist cannot give a coherent account of the experiences occurring in the example. For our purposes, we can think of it as a possible case where a subject is radically mistaken about the nature of their experience, in a way that seems to conflict with skepticism-proof access, and the functional constraints this (arguably) entails.

Shoemaker’s original illustrative example involves individuals called Bertha and Doris that have normal human color processing, but who each also have a “backup system” which functions just like the human color system, but which has a different physical realization.¹⁵ The backup system kicks in instead of the normal system when the subject needs it, and the subject notices no change, as it functions just like the normal system. To oversimplify a little bit, the problem the example creates is this: by a functional criterion for phenomenal similarity, the experience Bertha has when she sees a red thing with her normal visual system is qualitatively just like her backup-system red experience, because she cannot discriminate between them. But by a realization criterion, they are qualitatively different: we can set things up so that at least one subject’s backup system involves realizations of color experience that *must*, on a RR functionalist view, involve different experiences in response to color stimuli (in particular, we get this if Bertha’s backup system involves an inverted realization relative to Doris’s). So for at least one of Bertha and Doris, the functional criterion and the realization for phenomenal

14. Note that you could hold this role-realizer view without believing that phenomenal similarity has an *a priori* functional analysis (Shoemaker’s position). A similarity functionalist of the kind we are interested in here could have a more moderate version on which the similarity relations between experiences involve *a priori* functional constraints, but can’t necessarily be given a full *a priori* functional analysis. The problem I want to look at applies just as much to this weaker view.

15. Chalmers’s (1995, ch. 7) discussion of “dancing qualia” and “absent qualia” involves a similar example.

sameness give conflicting results: the role/realizer functionalist view is revealed to be incoherent.

The abstract reason for the trouble is this: the functional relations between realizers that determine similarity of experiences appear to be *external* relations. But then there is no guarantee that experiences realized by different tokens of the same type of realizer (which therefore [according to the argument] have the same intrinsic phenomenal nature) will *seem* the same to their subject: they could be wired together to seem phenomenally very different (or phenomenally different experiences could be wired together to seem exactly the same). So the similarity relations between experiences don't necessarily constrain their functional relations in the way that the similarity functionalist says that they should.

Shoemaker's example is unnecessarily complicated, in my view, and also involves an unnecessary feature that diminishes its force. In his example, Bertha and Doris are unable to distinguish normal from backup states in the sense that their *diachronic* comparisons do not distinguish them. But it's not clear how problematic this is: doesn't our ordinary conception of phenomenology allow that memory could play tricks on us, and make it seem like experiences at different times are the same, even though they are not?

In fact, this is not an essential feature of examples like this, nor is the complex structure of Shoemaker's case (which I did not fully describe). Here's a case inspired by Shoemaker's Bertha that I think makes the point more effectively. Consider a subject (I'll call her "Synchronic Bertha") who has two color visual systems rather than one, V1 and V2, which she can use *at the same time* (for example, perhaps she has eyes in the front and back of her head, connected to different systems). V1 involves a normal physical realization of color space, and V2 involves an inverted physical realization. So the type of realizer state that gets tokened when Bertha sees a green thing with V1 is the same as the type of realizer state that gets tokened when she sees a red thing with V2, and so on. So by a realization criterion, the qualitative character of a green-caused experience in V1 is very different from the qualitative character of a green-caused experience in V2. However, we can easily imagine setting things up so that she *does not notice this difference*. V1 and V2 might be wired together to enable Bertha to make "inter-system" color comparisons—that is, comparisons of the properties of external surfaces—across visual systems. And there appears to be no reason why, despite the inverted realization, these comparisons could not be made accurately. If we make the "transparency" assumption mentioned earlier (see footnote 10—this point deserves more extended discussion), that subjects compare color experiences simply by making externally directed color discriminations, then it will seem to Bertha as if her V1 and V2 color experiences are qualitatively alike.

This suggests that the following is possible: Bertha could be viewing a green object with one visual system and simultaneously viewing a different green object with the other visual system. These experiences might be indistinguishable from each other. Yet, by the realization criterion, they must be qualitatively very different, because her different systems involve inverted realizations of color experience.

So, by the realization criterion, Bertha is simultaneously having experiences of very different colors that nonetheless *seem* the same to her (that is: she is disposed to *judge* that they are the same; and insofar as the system responsible for “inter-system” color comparisons delivers an *experience* as of the relationship between the colors, she *experiences* them as being the same color).

This is exactly the kind of large error about experience that the anti-skeptical argument claims to be inconceivable. Imagine having an experience as of a large red circle against a black background in your left visual field and as of a similar green circle against a black background in your right visual field. Could it be that you are nonetheless disposed (even when attending scrupulously and alertly, etc.) to judge that they are the same shade (say, the same shade of green)? Again, this is exactly the kind of thing that the anti-skeptical argument says is inconceivable.

This might not seem so problematic for the anti-skeptical rationalist, because it might appear that they can easily just reject one of the assumptions that lead to the apparent possibility of a case like this. Actually I think there are four assumptions that lead to this kind of problem (although we may still have a problem without (4), as I explain below):

- (1) Realization Physicalism: Every instantiation of a phenomenal property is grounded in a (perhaps non-unique) set of physical events, its “realizer” (I’ll say more about what a “realizer” is immediately below).
- (2) Anti-Structuralism: The identities of phenomenal properties in a quality space is not fully determined by the similarity structure of the quality space.
- (3) Similarity Externalism: Phenomenal Similarity is determined by functional relations between the realizers of experiences, and these are external relations between the realizers.
- (4) Multiple Realizability: The same kind of experience can be multiply realized within a single cognitive system.

Rejecting (1) amounts to rejecting a fairly weak form of physicalism; below I’ll consider what the situation looks like from that perspective. Rejecting (2) involves embracing the kind of structuralist view already considered above. (Note that if we accept (1) and (2), and think that similarity functionalism is true, then we are likely to embrace role-realizer functionalism.) Prima facie, then, the similarity functionalist’s best bet is to reject (3) or (4).

The proposal Shoemaker himself tries to develop in response to this kind of problem (see Shoemaker 2006) is a view that I call “holistic role/realizer functionalism,” which I see as equivalent to rejecting (3) and accepting (1), (2), and (4). The proposal is best understood in terms of an ambiguity in the notion of “realizer.” A realizer could be a *total realizer*: a complete physical ground for an experience. But it could also be a *core realizer*: a more localized set of events that are directly relevant to what *kind* of experience is happening, but are not strictly sufficient for that kind of experience, or for any experience at all (Shoemaker 1981b). In the Synchronic Bertha example, these would be the kinds of physical events internal to one color system that are characteristic of one kind of color experience rather than

another, and which stand in the functional relations that determine phenomenal similarity.

The thing to notice is that (3) (Similarity Externalism) is not coherent if we take “realizer” to mean “total realizer.” Since similarity relations are internal relations between experiences, the total realizers for a set of experiences must be sufficient to determine these similarity relations, and therefore these relations must be grounded in *internal* relations between the realizers, not external functional relations. If, on the other hand, the realizers mentioned in (3) are *core* realizers, it’s perfectly coherent to take the functional relations between realizers to be external relations. This suggests that the Bertha problem only arises if we take the structurally transcendent qualitative character of an experience to be determined by features of its core realizer alone. Shoemaker’s (2006) idea is that rejecting structuralism doesn’t require us to hold this view. Instead, we could have a view on which this qualitative character is determined more holistically by other aspects of the total realizer also, in a way that avoids the problem generated by (3).

What would this look like? Shoemaker’s idea is that we need not take the role and realizer relevant to, say, a “red” experience as defined locally to one system, but rather as including all the local core realizers of this experience type in different parts of the system picked out by the way they are functionally connected (so these functional connections are now *internal* to the realization). For example, in my two-system color case, we need not take the role and realizer of a “red” experience as defined internally to one color system. Instead, we should think of the realizer more holistically, as involving the local “red” realizer in system V1, and also *the way it is connected to local “red” realizers in other visual systems*, local realizers of experiences that are indistinguishable from V1 red experiences. So, in Shoemaker’s Doris and Bertha example, we can say they have different experiences realized in their regular visual system, despite the fact this system is exactly the same in both of them: *the experiences are different because the regular system is connected to the backup system in a different way.*

On this view, Synchronic Bertha has color experiences that are qualitatively totally different from a normal human’s, despite having visual systems very similar to ours. Our red experiences are realized, let us say, by property R1, and our green experiences by property G1. But Bertha’s red experiences are realized jointly by R1 in V1, and G1 in V2, which we are to think of as a single holistically defined realization, giving rise to a different kind of experience.

This is an ingenious solution to the puzzle, but unfortunately it has serious difficulties as well. For one, there is the following problem that Shoemaker (2006) himself raised. What if V2 is rarely ever used? Then most of the time, the subject, using just V1 for color vision, has color vision that operates pretty much exactly like ordinary human color vision, including in its neural realization. Nonetheless, the holistic functionalist must say that this is not enough for it to be qualitatively like human vision, because the mere *potential* for “inter-system” comparisons involving V2 means that color experience is not realized in the same way here. This is already an implausible result. It gets even worse though,

when we imagine V2 gradually “atrophying” in such a way that is increasingly difficult for it to operate, and to communicate with V1. Eventually it becomes completely inoperative. The holistic functionalist must say that the end result of this process is that the subject gets color qualia just like a normal human’s. But when did the qualia shift? And if it did shift, presumably this happened in a way that the subject could not notice, a result that is itself in tension with Shoemaker’s functional criterion of phenomenal sameness (in its, admittedly problematic, diachronic version).

I think this already makes holistic functionalism problematic. But actually there is a much more serious objection, which Shoemaker does not appear to have noticed. I call this the *symmetry objection*. It is actually pretty salient in my two systems example: you may have already noticed the problem. It arises from the fact that the way V1 is hooked up to V2 is completely symmetrical to the way V2 is hooked up to V1 (see fig. 1). Suppose a red experience is realized locally in V1 by R1. The (total) realization of this red experience, on holistic functionalism, also includes the way that R1 is connected to states in V2 and other systems (how it is “globally connected”). In particular, R1 in V1 is wired up to be indistinguishable from the state realized by G1 in V2. Now, as we set things up, there are also experiences locally realized by R1 in V2. And because of symmetry, the way R1 is globally connected from V2 exactly matches the way R1 is globally connected from V1. So it looks like on holistic functionalism, R1 in V2 should realize the same kind of experience as it does in V1. But that’s the wrong result. V2 has an inverted realization relative to V1, so R1 experiences in V2 are actually indistinguishable from green experiences in system 1, not red experiences. So actually *holistic functionalism leads to exactly the same kind of incoherence that less holistic functionalist views (on which (3) is true) lead us to.*

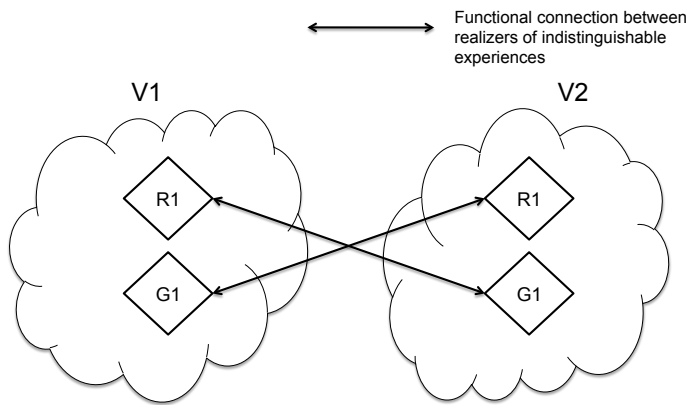


Figure 1. The experience core realized by R1 in V1 has the same holistic realizer as the experience realized by R1 in V2, and so Shoemaker’s holistic RR functionalist view entails they have the same phenomenology, but also that they have different phenomenology, given their functional connections.

Holistic functionalism is therefore not a good response to the Bertha problem. What about responses that reject assumption 4, the “multiple realizability assumption”? This is the assumption that conscious experiences of a particular external feature (such as the color property of a surface, understood as a mind-independent stimulus feature) can be multiply realized in a single cognitive system in a way that allows the resulting representations to be cross-compared. Rejecting this assumption gives you what I call the “common code” view: on this view, (1) different representations of the same feature can only be compared if they are translated into a centralized “common code,” and (2) experience of the feature is always realized by the common code representation and not by the earlier representations that cause the common code representations. In this way, it is not possible for there to be a Bertha-like scenario.

The problem with the common-code view is that it is both empirically and conceptually dubious. Conceptually dubious, in that it is unclear why intersystem (including intermodal) comparisons would necessarily require a common code, and even if they do, why conscious experience would have to be realized by it, rather than earlier in processing. Empirically dubious, in that it is not clear empirically that such comparisons in fact always do involve a common code (for example, there is evidence that distinct mechanisms are involved in millisecond timing for different modalities or even different spatial regions within a modality; subjects can nonetheless make temporal comparisons across modalities or regions [Mauk and Buanomano 2004; Johnston 2006]). Having said this, these objections to common coding are not quite as damning as those that hold for other views, so I see this response as worthy of further attention.

A final response is to reject the idea that experiences have any kind of physical ground at all—a fortiori they don’t have core or total realizers (i.e., to reject (1)). Of course, whether you like this response will depend on your antecedent views about the viability of physicalist or anti-physicalist views of various kinds. But more saliently, it’s very unclear how this actually helps, because the argument looks pretty good if we treat the “realizers” of experience as merely the nomic bases for experiences of different types. That is, even an anti-physicalist may want psycho-physical laws that (1) require a functional basis for the similarity structure of a quality space, guaranteeing that its subject cannot be subject to systematic introspective illusions, and (2) take the physical basis of experiences in a quality space to involve more than this functional structure, on the grounds that there can be different quality spaces with the same similarity structure, which must therefore have different physical bases but identical functional bases. But these assumptions create just the same problems we have been discussing, even though “physical basis” now is merely a nomic basis.

What all this suggests is that, contrary to similarity functionalism (and hence contrary to moderate rationalism), there are serious objections to views on which phenomenal similarity is functionally defined or functionally constrained in the way suggested by the anti-skeptical argument. Structuralism is false because it treats phenomenal properties too abstractly: in particular, it can’t handle the existence

of structurally identical yet qualitatively distinct quality spaces. But the alternative role-realizer similarity functionalist view is revealed to be incoherent on both a local and more holistic reading. The incoherence may depend on a multiple realizability assumption, but the assumption is both empirically and philosophically plausible. This is the argument against supposing experience to have features that satisfy anti-skeptical functional constraints.

V. THE ANTI-FUNCTIONALIST VIEW

Where does this leave us? It's important to see that the argument just given only requires rejecting the kind of functional constraints on phenomenal similarity that were suggested by the anti-skeptical argument; it doesn't require embracing a thorough-going anti-functionalism, on which the nature of experience is completely independent of functional structure in the brain. Such *strong anti-functionalism* is an extreme view; it's not even clear that there exist *any* physical properties that don't constrain the functional structure of the world; or at least, the kind of high-level properties that are constitutively involved in consciousness plausibly have at least a *partially* functional nature, or impose at least *some* functional constraints. The point is that if there are functional constraints on phenomenal properties, they allow for subjects to be radically mistaken about the similarity relations between their experiences.¹⁶ Call the view that rejects such functional constraints on phenomenal similarity but which *also* reject strong anti-functionalism, *moderate anti-functionalism*. We can use the term "anti-functionalism" to cover both moderate and strong versions.

It's helpful to describe two particularly important forms of anti-functionalism. One is an *externalist* view on which the phenomenal character of experience is determined entirely by the external stimulus property that the experience is an experience *of*, so that similarity relations between experiences are determined by objective similarity relations between external stimulus properties.¹⁷ This is an anti-functionalism view, because there is no guarantee that the subject's similarity judgments will align with the objective similarity relations between external stimulus properties which ground phenomenal similarity (don't be tempted to make the move of saying that internal functional structure constrains *which* external properties are experienced in such a way as to guarantee introspective reliability [see Byrne and Hilbert 2008]); that is to revert to a view on which the ultimate

16. One attractive view is that phenomenal similarity relations impose functional structure on the states inside individual perceptual channels, but similarity relations between experiences realized in different perceptual channels don't entail functional relations between those channels, such as cognitive dispositions to make similarity judgments (thanks to Ned Block for discussion here).

17. For some examples of externalist views, see Byrne and Hilbert 2008, Campbell 2002, Hill 2009, Lycan 2001, Tye 2000.

ground for phenomenal similarity is internal, not external). The other important form of anti-functionalism is an *internalist* version on which phenomenal character is determined by the features of internal neuro-functional realizer states; but as with external stimulus properties, the subject need not be wired up to have reliable knowledge of the similarity relations between these internal states (and hence the similarity relations between experiences).

What does anti-functionalism tell us about the existence of phenomenal properties? Our moderate rationalist thinks that there are anti-skeptical functional constraints on phenomenal properties; the anti-functionalist denies that experience has properties with any such features. So, it looks like the moderate rationalist's response at this point should be to conclude that if anti-functionalism is true, then *there are no phenomenal properties*. The alternative is to reject rationalism: we could take a more "diet" empiricist line, and say that anti-skeptical access is *not* part of the concept of phenomenology. Or there is an intermediate revisionary "diet" empiricist line, on which we should weaken the concept to exclude this connotation. On either of these latter views, it ends up being perfectly conceivable that we are systematically deluded about our experience: in that sense, introspection is just like external perception. This is a surprising result, but perhaps one we might be forced to accept.¹⁸

The situation is actually a little more complex than this, however, because of the distinction between ideal and absolute constraints mentioned earlier. A rationalist could interpret their constraints as idealized, and argue that Bertha is in an unusual, nonideal situation. How might that be? The most obvious kind of idealization here is to an assumption of rationality: it might be said that a priori constraints on phenomenal properties are primarily *epistemic*, and only have functional implications relative to rationality assumptions. Bertha is making mistakes about phenomenal similarity because she's not fully rational,¹⁹ not because the properties of her experience don't satisfy the relevant (idealized) functional constraints.

But this simply isn't a plausible construal of her predicament. We might concede that if someone is irrational enough, they might form massively erroneous beliefs about their experience. But there is no reason to think that Bertha is in any

18. Shoemaker's paradox appears to show that *all* a priori functional constraints on phenomenal similarity fail, so this might be taken to support a kind of empiricism on which there are no functional connotations to phenomenal concepts. But actually that is too quick: functional constraints could play a *reference fixing role*, even if they don't individuate phenomenal properties. For example, phenomenal properties could be picked out as those properties that introspective judgments *in fact, actually, or typically*, track, even if they don't have to track them. This is an interesting view: but note that it has the consequence that it is conceivable that our phenomenal concepts actually fail to refer to anything. Not only is a Bertha-like situation conceivable as a counterfactual scenario on this view; for all we know, we are actually in a Bertha-like situation, and so the reference-fixing role associated with a phenomenal concept might not actually be satisfied.

19. See Pautz (2011, p. 427, n. 22) for an example of such a view. A theorist like Pautz could just insist that a disposition to form perceptual beliefs that reflect experience in a certain way is a condition on rationality, but that isn't an independently plausible reading of the Bertha case, I don't think.

way irrational. She makes mistakes (if that is what they are), just because of the way her *perceptual* systems are linked up—not because of anything about how her *cognitive* system is set up, or any kind of defective or nonideal functioning; this makes it very strange to say that this is a rational defect. She might be perfectly rational and still be prone to these errors.

Another option here (suggested to me by David Chalmers; see also Chalmers 1996, ch. 5) is to say that the anti-skeptical intuition can be accommodated by the view that having an experience of a certain type gives a subject an automatic justification for believing they have this type of experience (call this the “justification constraint”). So a subject can’t be in massive error about their experience, in the sense that if they form the beliefs about experience that they are justified in forming, then they will not be mistaken (by contrast, we might want to allow that perceptual beliefs can be based on all available justification, but still radically mistaken). The problem with subjects like Bertha (or *us* in many cases also, where we don’t, or cannot, access aspects of our experience) is not that they are irrational, but simply that they are not in a position to take advantage of the justification they have for an accurate belief about their experience.

But the justification constraint, as appealed in this context, seems both too weak and too strong, in different ways. It seems too weak in that it’s not clear that it really captures the strength of the anti-skeptical intuition. Suppose I assume that right now I am completely rational, alert, and that my perceptual/introspective capacities are working exactly as they are designed to work. If all we have in place is the justification constraint, then it is still conceivable that I am massively radically mistaken about the nature of my current experience. I might be in a Bertha-like predicament, where, despite being in conditions that are in many ways epistemically ideal, I am unable to access the justification I have for accurate introspective beliefs, and am instead in massive error.

The constraint also seems too strong, in that it appeals to the dubious idea of a justification for a belief that a subject may be completely incapable of accessing. There may be cases where we want to allow for this kind of thing: for example, a subject might believe the premises of an argument that could enable them to derive a certain conclusion, but be unable to derive the conclusion because of performance limitations. But Bertha’s situation doesn’t intuitively involve a performance limitation in this way; her cognitive system might be working optimally as it was designed to work.

To get at it another way: the suggested lesson of the Bertha case is that our access to phenomenal similarity is quasi-perceptual in that it relies on a fallible causal mechanism that could in principle yield systematically deceptive results. In that way, our access to phenomenal similarity is like our perceptual access to external world facts. One could analogously respond to external world skepticism by being more “externalist” about our perceptual evidence, including external world facts themselves as part of our basic evidence; in that way, it might be said that we have evidence definitively ruling out skepticism. But even if we make this externalizing move, surely we would not include external facts that the subject has no

access to. Similarly, it would be implausible to treat as basic evidence phenomenal similarity facts that the subject has no access to.

Suppose this point is conceded. A fan of the justification constraint could still modify their view as follows. They could say: even if massive systematic error is possible, and even if Bertha doesn't have justification for her similarity beliefs, *we* do, because we *do* have access to the phenomenal similarity facts, which can therefore be taken as part of our basic evidence. Moreover, since these similarity facts are part of our basic evidence, we have definitive reason for ruling out phenomenal skeptical scenarios, even if Bertha doesn't; in that sense, they are inconceivable.²⁰

However, this is still unsatisfactory. First, it generates a second-order skeptical challenge: how do I know that these similarity relations are part of my basic evidence? One need not believe in a dubious general principle to the effect that justification requires second-order knowledge of the source of justification to find this particular second-order challenge compelling. Once you see that you have access to phenomenal similarity only through a fallible causal mechanism that could yield Bertha-like mistakes, it is surely reasonable to wonder whether you could be subject to such mistakes yourself (compare again how we intuitively react to externalist responses to external-world skepticism that claim that facts about the external world are part of our basic evidence, and so skeptical scenarios are in a certain sense inconceivable).

Second, and more significantly, the response misconstrues the intuition behind the original anti-skeptical argument. We can construe the argument as saying that not only is it inconceivable that *I* am in a skeptical scenario, but also it is inconceivable that *anyone* who has an experience like mine (and who is not completely irrational or distracted) could be completely wrong about their experience. In conceding that the Bertha scenario is possible, the response gives this up, and therefore does not really do justice to the anti-skeptical intuition (or the resulting rationalist idea that there are functional/epistemic constraints on phenomenal similarity). At best it gives us a weak surrogate sense in which we are able to rule out phenomenal skepticism, one that doesn't entail an interesting form of rationalism.

To sum up, we have a dilemma here: if we try to prevent Bertha from being a counterexample to similarity functionalism by weakening our functional constraint through idealization, then we risk ending up with something far too weak to capture the anti-skeptical intuition, or otherwise implausible. But then it looks like it might be that the only way of doing justice to the intuition makes Bertha into a counterexample.

What if, instead of trying to save similarity functionalism in this way, we react to Bertha with a form of empiricism? We could deny that there ever was such

20. To put it another way: the Bertha-generated concern, it might be said, is that it is consistent with how things seem to you that you are totally wrong about your experience. But "how things seem to" includes the similarity relations between experience you *in fact* have introspective access to, and for you (even if not for Bertha), they definitively rule out your being in a Bertha-like scenario.

Compare the responses of Hill (1991, ch. 6), Chalmers (1996, ch. 5), Levine (2001, ch. 5) to Shoemaker's zombie argument.

an anti-skeptical epistemic constraint on experience; or we could argue that the concept of phenomenology ought to be weakened to remove any such constraint (such a revisionary move of course entails that the rationalist was correct about the original concept). I suspect this will seem like an attractive option to many, and I'm not opposed to it. But I think it is important to understand where it leaves us.

If phenomenal skepticism is conceivable, then our epistemic situation vis-à-vis experience is much more similar to our situation vis-à-vis our perceptual connection to the external environment than we might have initially thought. The irony, as suggested at the beginning of the paper, is that although removing special access constraints on phenomenal properties makes their existence in one way safer, it also makes their existence and nature in another way far more insecure: we now have to admit that it's conceivable that we are completely wrong about them. In fact the situation could be even worse; it's not clear that the conceivability of a Bertha-like predicament doesn't slide us in the direction of it being conceivable that none of the apparent phenomenal properties of experience exists at all, or even that experience itself doesn't exist. As mentioned earlier, it is actually quite plausible that phenomenal reference is subject to *contrast constraint*, much as external reference appears to be. But if so, then the truth of phenomenal skepticism might make it impossible to even conceptualize phenomenal aspects of experience; so it is conceivable that our attempts to think about the phenomenal properties of experience may be completely failing: in that sense, *the phenomenology of experience might not exist*.

This troubling epistemic situation for the empiricist is compounded by another, related, worry (briefly mentioned earlier), that I think deserves closer attention. The empiricist's strategy of completely purging phenomenal notions of theoretical content risks making them unacceptably indeterminate or even devoid of content, inviting a Frankish-style conceptual eliminativism. The problem is particularly serious if we assume that materialism is true; then there are a plethora of complex high-level properties of brain states that are candidates for being the semantic values of phenomenal concepts. The empiricist is suggesting that we can directly pick one of these properties out in a pure way, using a completely theoretically neutral phenomenal concept. But that just seems like magic—what about the concept could possibly select one, or even a small range, out of the many complex properties in the vicinity? I think the radical empiricist needs to urgently address this question (one option is simply to bite the bullet and accept that phenomenal concepts are highly indeterminate [Papineau 2002, ch. 7]).

If such concerns have force, they actually suggest an interesting positive constructive project: to find a conception of phenomenology that has enough theoretical content to determinately pick out properties of experience, but not involve mistaken presuppositions about experience. If we are in this situation, that's of course a far cry from a phenomenology being the kind of thing that we can just easily and effortlessly pick out; but I do see it as a predicament we might turn out to be in.

To sum up—if we reject structuralism, the possibility of Bertha-like scenarios appears to show that there aren't functional constraints on phenomenal similarity,

and so phenomenal skeptical scenarios are possible (even if they are hard to conceive of first personally). This means that a rationalist either has to accept that phenomenal properties don't exist, or switch to a more (possibly revisionary) empiricist position. I haven't been able to fully assess the merits/problems with such an empiricist view here, but I have suggested that rejecting anti-skeptical access as a constraint at least has the interesting consequence that our epistemic access to phenomenal properties is less secure than we might have thought—indeed it may be conceivable that the phenomenal properties that your current experience appears to have don't exist.

Finally, there are two issues for further discussion I want to flag. First, if phenomenal skeptical scenarios are conceivable, could I get evidence that I'm actually in one? For example, could I find out that I'm wired up in a Bertha-like way, and so get third-personal evidence that my introspective judgments are systematically mistaken?²¹ I must confess that I'm not sure I could ever accept such evidence—I just couldn't shift my Cartesian conviction that experience is more or less how I believe it to be. But what to make of this? Is it that I can't really accept phenomenal empiricism, or that I'm being irrational? Or can introspective evidence always trump third-personal evidence, even if anti-skeptical rationalism is false? This deserves more discussion.

Second, if we do not have special access to phenomenal aspects of our experience, why does it seem to us like we do, or even inconceivable that we don't? I will have to leave detailed discussion of this for another occasion, but briefly: I'm tempted to think that it is an example of how imagination misleads us in our understanding of experience. We cannot imagine experience simultaneously having a certain similarity structure, without imagining having *access* to this similarity structure; we mistakenly infer that experience and access always go together. There may be other examples where imagination misleads us in this way. For example, Prinz (2012, ch. 8) argues that the belief that parts of experience are necessarily always “unified” (Bayne 2010) derives from the facts that experiences we jointly attend to or imagine having always are unified. Another example is the determinateness of experience: it seems inconceivable that there are creatures that only indeterminately have consciousness—either the lights are on or they aren't! But again, it is plausible to say that we are misled here by the fact that we can't imagine indeterminate consciousness from the inside; that may reflect the nature of imagination, and not the nature of consciousness, however.

21. For a realistic example, consider the phenomenology of duration perception (in particular, of short durations [less than 2 seconds] that can be directly perceived). Auditory processing has much higher temporal resolution than visual perception, and so is in one respect much *faster*. Suppose someone took this as a reason for saying that visual and auditory temporal phenomenology are very different—in particular, time is “flowing” much more slowly in audition (note that the analogous inference for intersubjective comparisons of temporal experience across different organisms is extremely tempting). This is very hard to accept, because we are able to make intermodal introspective comparisons with respect to duration phenomenology, and we don't notice any such discrepancy. But what if this is a systematic introspective illusion?

VI. CONCLUSION

This paper had a rather long and complex structure, so let me end by summing up where we have got to. I distinguished rationalist and empiricist views of phenomenal concepts, and set out to assess the prospects of a moderate rationalist view. I suggested that a modified version of Shoemaker's anti-skeptical argument for functionalism supports the rationalist view that we conceptualize phenomenal properties as having a special epistemic status—in particular, we cannot conceive of certain kinds of skeptical scenarios about experience obtaining. This supports the claim that we conceive of the phenomenal similarity structure of experience as subject to functional constraints—it is accessible to us in a way that rules out systematic error.

But although I think it's plausible that we *conceive* of phenomenal similarity structure as accessible in this way, I raised some objections to the idea that experience actually *is* like this. In particular, I argued that Shoemaker's paradox suggests that radical systematic error about the similarity structure of experience is after all possible, even if hard to conceive of first personally, leading to the conclusions just mentioned above: experience may have "qualia" in a thinned-down sense, but if so, we can't be sure that it does.

ACKNOWLEDGMENTS

Thanks to Ned Block, David Chalmers, Chris Hill, and Adam Pautz for helpful comments and discussion.

REFERENCES

- Baars, B. J. 1997. *In the Theater of Consciousness: The Workspace of the Mind*. Oxford: Oxford University Press.
- Balog, K. 1999. "Conceivability, Possibility, and the Mind-Body Problem." *Philosophical Review*, 108 (4):497–528.
- Bayne, T. 2010. *The Unity of Consciousness*. Oxford: Oxford University Press.
- Block, N. 1980. "Are Absent Qualia Impossible?" *Philosophical Review*, 89: 257–74.
- Block, N. 1995. "On a Confusion about a Function of Consciousness." *Behavioral and Brain Sciences* 18: 227–47.
- Block, N. 2007. "Consciousness, Accessibility, and the Mesh between Psychology and Neuroscience." *Behavioral and Brain Sciences* 30(5–6): 481–99.
- Byrne, A., and D. Hilbert. 2008. "Basic Sensible Qualities and the Structure of Appearance." *Philosophical Issues* 18(1): 385–405.
- Campbell, J. 2002. *Reference and Consciousness*. Oxford: Oxford University Press.
- Chalmers, D. J. 1998. "On the Search for the Neural Correlate of Consciousness." In *Toward a Science of Consciousness II: The Second Tucson Discussions and Debates*, vol. 2, ed. S. R. Hameroff, A. W. Kaszniak, and A. Scott, 2, 219. Cambridge, MA: MIT Press.
- Chalmers, D. J. 2003. "The Content and Epistemology of Phenomenal Belief." In *Consciousness: New Philosophical Perspectives*, ed. Q. Smith and A. Jokic, 220–72. Oxford: Oxford University Press.

- Chalmers, D. J. 1995. "The Conscious Mind: In Search of a Fundamental Theory." Oxford: Oxford University Press.
- Clark, A. 1992. *Sensory Qualities*. New York: Clarendon.
- Dennett, D. C. 1988. "Quining Qualia." In *Consciousness in Contemporary Science*, ed. A. J. Marcel and E. Bisiach, 42–77. Oxford: Oxford University Press.
- Frankish, K. 2012. "Quining Diet Qualia." *Consciousness and Cognition* 21(2): 667–76.
- Gescheider, G. A. 1997. *Psychophysics: The Fundamentals*. Mahwah, New Jersey: Psychology Press.
- Goodman, N. 1951. *The Structure of Appearance*. Cambridge, MA: Harvard University Press.
- Harman, G. 1990. "The Intrinsic Quality of Experience." *Philosophical Perspectives* 4: 3152.
- Hill, C. 1991. *Sensations: A Defense of Type Materialism*. New York: Cambridge University Press.
- Hill, C. 2009. *Consciousness*. Cambridge: Cambridge University Press.
- Hinton, J. M. 1967. "Visual Experiences." *Mind*, 217–27.
- Johnston, A., D. H. Arnold, and S. Nishida. 2006. "Spatially Localized Distortions of Event Time." *Current Biology* 16(5): 472–79.
- Kind, A. 2001. "Qualia Realism." *Philosophical Studies* 104(2): 143–62.
- Laming, D. R. J. 1997. *The Measurement of Sensation*, vol. 30. Oxford: Oxford University Press.
- Lee, G. 2013. "Materialism and the Epistemic Significance of Consciousness." In *Current Controversies in Philosophy of Mind*, ed. U. Kriegel. London: Routledge.
- Levine, J. 2001. *Purple Haze: The Puzzle of Consciousness*. Oxford: Oxford University Press.
- Levine, J. 2004. "Thoughts on Sensory Representation: A Commentary on Austen Clark's Theory of Sentience." *Philosophical Psychology* 17(4): 541–51.
- Loar, B. 1990. "Phenomenal States." *Philosophical Perspectives* 4: 81–108.
- Lockhead, G. R. 1992. "Psychophysical Scaling: Judgments of Attributes or Objects?" *Behavioral and Brain Sciences* 15(3): 543–58.
- Lycan, W. G. 2001. "The Case for Phenomenal Externalism." *Noûs* 35(s15): 17–35.
- Martin, M. G. F. 2006. "On Being Alienated." In *Perceptual Experience*, ed. Tamar Gendler and John Hawthorne. Oxford: Oxford University Press.
- Mauk, M. D., and D. V. Buonomano. 2004. "The Neural Basis of Temporal Processing." *Annual Review of Neuroscience* 27: 3007–40.
- Papineau, D. 2002. *Thinking about Consciousness*. New York: Clarendon Press.
- Pautz, A. 2011. "Can Disjunctivists Explain Our Access to the Sensible World?" *Philosophical Issues* 21(1): 384–433.
- Prinz, J. 2012. *The Conscious Brain*. Oxford: Oxford University Press.
- Quine, W. V. O. 1980. "Two Dogmas of Empiricism." In *From a Logical Point of View: 9 Logico-philosophical Essays*. Cambridge, MA: Harvard University Press.
- Rey, G. 1983. "A Reason for Doubting the Existence of Consciousness." In *Consciousness and Self-Regulation*, vol. 3, ed. R. Davidson, G. E. Schwartz, and D. Shapiro. New York: Plenum Press.
- Rosenthal, D. 2010. "How to Think about Mental Qualities." *Philosophical Issues* 20(1): 368–93.
- Schwitzgebel, E. 2008. "The Unreliability of Naive Introspection." *Philosophical Review* 117(2): 245–73.
- Schwitzgebel, E. 2011. *Perplexities of Consciousness*. Cambridge, MA: MIT Press.
- Shepard, R. N. 1978. "On the Status of 'Direct' Psychophysical Measurement." In *Minnesota Studies in the Philosophy of Science*, vol. 9, ed. C. W. Savage. Minneapolis: University of Minnesota Press.
- Shoemaker, S. 1975. "Functionalism and Qualia." *Philosophical Studies* 27(5): 291–315.
- Shoemaker, S. 1981a. "Absent Qualia Are Impossible—a Reply to Block." *Philosophical Review* 90: 581–99.
- Shoemaker, S. 1981b. "Some Varieties of Functionalism." *Philosophical Topics* 12(1): 93–119.
- Shoemaker, S. 1982. "The Inverted Spectrum." *Journal of Philosophy* 7: 357–81.
- Shoemaker, S. 2006. "The Frege-Schlick View." In *Content and Modality: Themes from the Philosophy of Robert Stalnaker*, ed. J. Thomson and A. Byrne, 18–33. Oxford: Oxford University Press.
- Stalnaker, R. 1999. "Comparing Qualia across Persons." *Philosophical Topics* 26: 385–405.
- Tye, M. 2000. *Consciousness, Color and Content*. Cambridge, MA: MIT Press.
- Tye, M. 2002. "Representationalism and the Transparency of Experience." *Noûs* 36(1): 137–51.