

Forthcoming in Classic Arguments: The Knowledge Argument, ed. Sam Coleman, Cambridge University Press

THERE'S NOTHING ABOUT MARY

David Rosenthal
Philosophy and Cognitive Science
Graduate Center, City University of New York

Abstract: Mary is shown something and told it's red, and she says, "This is what it's like to see red." On Jackson's (1986) knowledge argument, what she says expresses factual knowledge that the experience is a case of seeing red. In addition, Jackson argues, Mary could not have had that factual knowledge before first seeing something red; though Mary's books cover all facts about the physical makeup of the world, the knowledge her remark expresses could not have been in her books.

Some have argued that the knowledge is both factual and new to Mary by urging that the knowledge is cast in part in phenomenal concepts, concepts one cannot have without having had the relevant qualitative experience. The nature of phenomenal concepts ensures that the knowledge is new, and because they're concepts the knowledge is factual. But positing phenomenal concepts faces challenges that likely cannot be met.

Others have argued that though the knowledge is new to Mary, it's not factual. Some maintain that it's just knowing how to recognize or imagine a case of seeing red, others that it's just knowledge by acquaintance. I'll argue that both alternative types of knowing bring descriptive knowledge along with them. So unless one shows that this descriptive knowledge is not new to Mary, those accounts fail to sidestep Mary's having new factual knowledge.

Mary's new knowledge pertains to the qualitative character of her new experience. So to assess what knowledge Mary might have on first seeing red, we must understand what it is to be in a qualitative mental state. Without being clear about what qualitative character is, we cannot evaluate the claim that Mary has factual knowledge about it. There are two competing pictures of what conscious qualitative character is. On one Mary does gain new knowledge, but that the knowledge is not factual. On the other, which I argue is independently preferable, the knowledge Mary might have on first seeing red is indeed factual, but it won't be new to her; it will have been in her books. If Mary does have the factual knowledge needed for Jackson's argument, it is knowledge she would have before first seeing red.

Introductory

Frank Jackson's (1986) Mary is confined in a room in which the visual stimuli are all grayscale, so that her visual experiences have all been achromatic. Nonetheless, from books and television lectures, she has gotten all factual knowledge one can get from any source that pertains to the having of conscious visual experiences. It could indeed be "everything [factual] there is to know about the physical nature of the world," based on "completed" science. But, Jackson writes, "[i]t seems ... that Mary does not know all there is to know. For when she is let out of the black-and-white room or given a color television, she will learn what it is like to see something red, say" (1986, 291; emphasis Jackson's).

Jackson presents this as an argument against physicalism. Though Mary has "complete physical knowledge" (291) about the world before first seeing red, she nonetheless learns something new on being presented with a red stimulus. And that new knowledge is factual; it's the fact that this is what it's like to see red. Since Mary learns a new fact and already knew all the physical facts, the new fact must be nonphysical. This is the so-called knowledge argument against mind-body physicalism.¹

My focus here will not be on physicalism, though the implications for the anti-physicalist argument will be plain. Rather, my concern is how to understand the case Jackson describes and whether standard descriptions (e.g., in Ludlow, Nagasawa, and Stoljar 2004) are correct. Careful focus on that case, I'll argue, shows that though Mary is plainly in a mental state that is new to her on first seeing something red, the factual knowledge that this is what it's like to see red would not be new to her. It would have been contained in her books.

I. What Mary Learns

Jackson describes what Mary learns simply as "what it is like to see something red." What it's like to see something red is, for these purposes, the conscious qualitative character of the relevant state of seeing. So we can equivalently describe Mary as learning what the conscious qualitative character of seeing red is.

Both descriptions describe Mary's new knowledge not by way of a 'that' clause, but rather as coming to know 'wh'.² To describe somebody as knowing 'wh' is to describe their

¹ Cf. Jackson (1982). The basic argument is not new; "A blind man could know the whole of physics, but he could not know what things look like to people who can see, nor what is the difference between red and blue as seen" (Russell 1927, 182). I am grateful to Galen Strawson for this reference. See also Robinson (1982, 4-5).

² Cf. Daniel Stoljar's (2016) subtle and penetrating semantic account of 'what it's like' statements, on which, to a first approximation, 'There is something it is like to have a toothache' means that there is a way that having a toothache affects one. So knowing

knowledge in an incomplete way. If I describe you as knowing where the treasure is buried or what's behind the door, I'm saying that you know the answer to the question where the treasure is buried or what's behind the door. But knowing those answers is knowing that something is the case. My description in terms of knowing 'wh' abstracts from the knowledge 'that' I take you to have, knowledge that something is the case.³

Since knowing 'wh' is knowing the answer to the question expressed by the 'wh' clause and that the answer can be expressed by a 'that' clause, there is an important connection between knowing 'wh' and knowing 'that'. One knows 'wh' only if one knows that something relevant is the case (e.g., Vendler 1972). If one knows what's behind the door, one can in some way describe what's behind the door, even if in incomplete terms, and that would express knowledge that something is the case. Without some relevant knowledge that something is the case, there's nothing in virtue of which one knows 'wh'.

So when Mary knows what it's like to see something red or what the conscious qualitative character is that's characteristic of seeing red, she must have at least some knowledge that something is the case, something relevant to her knowledge 'wh'. And if her knowing 'wh' is new to her, that factual, descriptive knowledge must be as well.

But Jackson tells us nothing about it's what that new factual knowledge might be. What is it that Mary comes to know to be the case on first seeing something red? What is the new factual knowledge that underwrites her new knowledge 'wh'?

The natural candidate is that Mary comes to know that this is an example of what it's like to see red or of the conscious qualitative character of seeing red, where 'this'⁴ refers to her new experience. These seem to be clear cases of knowing that something is the case, which could underwrite Mary's coming to know 'wh'—what it's like to see red and what the qualitative character of seeing red is.

Mary's new knowledge is that whatever 'this' refers to is an example of seeing red. So the foregoing account of Mary's knowledge requires that we be able to tell which experience 'this' refers to independent of its being an example of seeing red. Mental pointing is unlikely to help here; without an independent account such an appeal simply labels the need to fix the reference of 'this'.

But mental pointing isn't needed. We can fix the reference of 'this' descriptively, by the location of the experience in Mary's visual field relative to other describable items also

what it's like to have a toothache is knowing what that way is (pp.1171-2). The 'wh' construal is preserved for knowing what it's like.

³ Some (e.g., Higginbotham 1996) have urged that knowing 'wh' is definable in terms of knowing that something is the case; but it is doubtful that actual definability is possible (e.g., Farkas 2015 and George 2013).

⁴ More precisely, the mental analogue of 'this', though I'll typically take that as understood.

present in that field. So one way to capture the factual knowledge Mary is supposed to gain on first seeing something red is that her new experience is an example of what it's like to see something red or, equivalently, an example of the conscious qualitative character of seeing red.⁵

But this way of capturing the factual knowledge Mary is supposed to gain does not, as Jackson notes, result in a problem for physicalism. Mary's experience of seeing red is by hypothesis new to her. So the knowledge simply that her new experience is an example of what it's like to see red is also new to her; Mary's books could not have covered facts about a token state of Mary's that had not yet occurred when those books were written. As Jackson puts it, "physicalist and nonphysicalist alike can agree" that Mary "could not have known facts about her [novel] experience of red, for there were no such facts to know" (1986, 292).

But that has no bearing on Jackson's anti-physicalist argument. The problem for physicalism, he urges, is not that on first seeing something red, Mary learns that her new experience is an example of what it's like to see something red. It's that she learns that

⁵ Indeed, it is arguable that all pointing and all demonstratives in thought or speech can be fully cashed out in descriptive terms. Zenon Pylyshyn (2007, ch. 2, esp. 2.1) has contested this, arguing that some mental demonstratives are ineliminable. His argument appeals to his ingenious and highly influential experimental work in multiple-object tracking, in which subjects can visually track roughly four moving targets when presented with four moving distractors. This, he writes, "make[s] it very unlikely that objects are tracked by regularly updating a description that uniquely picks out the objects" (17).

But that is not obvious. Subjects cannot consciously report descriptively how they track roughly four objects, but that does not show that the tracking is not achieved by descriptive information that either is unconscious or conscious for just a moment until no longer needed. Consider tracking just a single object with a single distractor, or perhaps two. In that case, most subjects could doubtless report in descriptive terms how they track. Adding objects to the limit of subjects' ability to track imposes substantial additional attentional load (e.g., Lavie 1995, 2010), diminishing other cognitive abilities, including conscious awareness (Lavie *et al* 2014). So the failure to report in descriptive terms the tracking of all four might well be just a casualty of attentional load, even if tracking is actually achieved by descriptive processing. One test might be to freeze the movement of tracked items from time to time and ask subjects to give whatever descriptive information about tracking they can, perhaps for just one or two items.

Subjects are of course visually acquainted with the items, though as argued in the next section such acquaintance does not figure in any operative conceptual content. And we have so far no reason, even in this ingenious experimental paradigm, to think that a irreducible demonstrative component figures in the tracking, as against purely descriptive information. This in turn casts serious doubt on any appeal to demonstratives in capturing the cognitive content of new factual knowledge Mary allegedly gains.

the property her new experience exemplifies is also what it's like for others to see something red. She learns that the conscious qualitative character exemplified by her new experience is exemplified generally by people's experiences on seeing something red.⁶

And others had seen red before Mary did, and her books would have described the neurophysiology and behavioral consequences of their doing so. But the books could not, according to Jackson, have included the fact that Mary's new conscious qualitative character is that of experiences that people generally have on seeing something red. So that's something new Mary learns that could not have been in her books.

This raises two questions. One is whether, on Mary's learning that her new experience exemplifies what it's like to see something red, she does also learn what it's like for others to see something red. The other is whether Mary's books could have contained descriptive knowledge of the conscious qualitative character of seeing something red, the qualitative property that characterizes others past experiences of seeing red. To answer these two questions we must get clear about what that conscious qualitative character consists in, that is, what the property is that Mary has knowledge about.

II. Phenomenal Concepts

To sustain Jackson's argument, Mary must on first seeing something red have knowledge that is both factual and new to her. We need to know what properties that new factual knowledge would be about. And a promising way to become clear about the properties a particular type of factual knowledge is about is by appeal to the conceptual content of that knowledge. And it has seemed to many that the relevant concepts here are so-called phenomenal concepts.

There are many accounts of phenomenal concepts (Alter and Walter, eds, 2007, and Nida-Rümelin 2009). But they have in common that one can have and deploy a particular phenomenal concept only if one has had the relevant qualitative experience. That's because the qualitative experience itself is the mode of presentation in virtue of which the phenomenal concept applies to a conscious qualitative state.

Accounts of phenomenal concepts vary in respect of how the conscious qualitative property serves as the mode of presentation in virtue of which phenomenal concepts apply

⁶ It's unclear that all discussions of the knowledge argument cast the problem in this way. Some seem to proceed on the assumption that the problem for physicalism might be simply learning what the conscious qualitative character is of Mary's new experience. But I'll follow Jackson's treatment on this.

In casting things this way Jackson must be tacitly assuming that undetectable inversion of conscious mental qualities cannot occur, since if it could, Mary could not know what it's like for others just by knowing what it's like for her. I won't pursue that issue here.

to those conscious qualitative properties. But I'll focus here on this pivotal common feature. Mary comes to know that this is what it's like to see something red, where the concept in her knowing that applies to the property of what it's like to see something red is a phenomenal concept.

If Mary's knowledge involves phenomenal concepts so construed, her knowledge is plainly new to her on first seeing something red. Knowledge cast in terms of phenomenal concepts could not have been in any books, since one cannot have phenomenal concepts at all without having the relevant qualitative experiences. We need know nothing more about conscious qualitative character to conclude that knowledge about it could not be in Mary's books. And since phenomenal concepts are concepts, knowledge cast in terms of them should be expressible by 'that' clauses. So Mary's knowledge should be genuinely factual, descriptive knowledge, as Jackson's argument requires.

One might have some concern about the appeal concepts so perfectly tailored to deliver the results Jackson's argument requires.⁷ Since phenomenal concepts are exactly what the argument calls for, we should want some reason to think such concepts actually occur, a reason that's independent of a desire to see Jackson's argument as sound. We cannot dispel this concern simply by providing an account of what phenomenal concepts are. We must have an independent reason to think there are any such concepts, and that they actually occur in one's thinking about one's own conscious qualitative states.

But set that methodological concern aside. There is a wholly distinct challenge about whether phenomenal concepts could be genuinely conceptual at all. The defining characteristic of phenomenal concepts is their tie to qualitative experiences. But the standard test for whether a mental item is conceptual is whether it has distinctively conceptual connections to other mental items that are themselves plainly conceptual.

If there are phenomenal concepts that occur in an individual's thinking or knowledge, they apply to the qualitative experiences of that individual, as they are conscious for that individual. So we could expect such phenomenal concepts to have distinctively conceptual ties with concepts that apply to qualitative experiences independent of whether one has had the experiences, concepts one could apply indifferently to oneself and to others. An example would be the concepts that a blind person could apply to the visual experiences of others.

As noted above, Jackson himself assumes that on coming to know that her new experience exemplifies the qualitative character of seeing red, Mary also comes to know that the experiences others have on seeing red have the same the qualitative character that her new experience exhibits. "[T]he knowledge Mary lacked which is of particular

⁷ A point of clarification: Phenomenal concepts are typically invoked to explain how Mary can gain new factual knowledge without thereby undermining mind-body physicalism. My concern here is only tangentially about mind-body physicalism; I am primarily concerned to contest that Mary does get new factual knowledge. And it is that claim that the positing of phenomenal concepts is so perfectly tailored to fit with Jackson's account.

point for the knowledge argument against physicalism is knowledge about the experiences of others, not about her own" (1986, 292; emphasis Jackson's). For that to be so, Mary's new knowledge must be cast in concepts that apply indifferently to herself and to others or, if not, then at least in concepts with robust conceptual ties to concepts that do apply indifferently to oneself and to others.

But phenomenal concepts cannot satisfy this condition. The mode of presentation in virtue of which phenomenal concepts apply to qualitative experiences ineliminably involve the very qualitative experiences to what they apply. So phenomenal concepts cannot themselves apply to others' qualitative experiences. If they did, one could have phenomenal concepts without having any relevant qualitative experiences, undermining the type of mode of presentation phenomenal concepts are said to have.

So the only way phenomenal concepts can function as concepts at all is for robust conceptual ties to hold between the phenomenal concepts, in terms of which Mary's knowledge is cast, and other, nonphenomenal concepts that apply to others' qualitative experiences. Such concepts would apply to others' qualitative experiences at least in part on the basis of third-person considerations.

But it's far from clear that any clearly conceptual connection can hold between a concept of seeing red that applies independently of whether one has seen anything red and a putative concept that applies only to oneself and only if one has seen something red, and indeed in virtue of one's having done so. Conceptual connections obtain in virtue of the application conditions of the relevant concepts, and the requirement that phenomenal concepts apply only to the individual that has the concepts makes connections with concepts that apply independent of that constraint unlikely. And the connection must be uncontroversially conceptual, since that's key to whether a mental item is conceptual. This is the main challenge for any appeal to phenomenal concepts.

Ordinary concepts operate descriptively, by applying to a tolerably well-defined range of items. So there's typically no difficulty in articulating distinctively conceptual connections among them. But an advocate of phenomenal concepts might contest the use of such connections to test whether phenomenal are genuinely conceptual. Phenomenal concepts are by hypothesis a special case.

But the knowledge argument requires that Mary's new knowledge be factual; so the knowledge must be expressible in conceptual terms. Without some way to show that phenomenal concepts are genuinely conceptual, we have no reason to hold that positing them can sustain the claim that Mary does gain new factual knowledge. And there seems no way to show that phenomenal concepts are genuinely conceptual except by appeal to conceptual connections with other concepts.

The issue about conceptual ties between phenomenal concepts and other concepts reflects a deeper difficulty that often arises in discussions of mental phenomena. Plainly we have first-person access to our own mental states, access independent of observing ourselves and independent of conscious inference. But we also have access to others'

mental states, and that third-person access does rest on observation and inference. And though our first-person access is only to our own states, the states of others to which we have third-person access are states of the very same sort as those to which we have first-person access. If I say I'm in pain and you deny that I am, we are not talking past one another; we are talking about states of the very same sort (see Rosenthal, 2018).

This gives us compelling reason to hold that whatever concepts apply to our own experiential states must have distinctively conceptual connections with concepts that apply indifferently to experiential states of ourselves and others. The problem with phenomenal concepts is the stricture that one cannot have them at all without having had the relevant experiential states. This makes it difficult to see how conceptual ties might hold between them and nonphenomenal concepts of experiential states, and so casts doubt on whether so-called phenomenal concepts are genuinely conceptual at all.

An advocate of phenomenal concepts might urge that they can, after all, apply to others' qualitative experiences. But it is unclear how that can fit with such concepts' having modes of presentation always inextricably tied qualitative experiences of one's own. And if phenomenal concepts aren't inextricably tied to one's own qualitative experiences, it's unclear how they can sustain Jackson's claim that Mary comes to have new factual knowledge. Simply claiming that one can extrapolate from the application of phenomenal concepts to oneself to applying them to others does nothing but ignore the challenge.

The most influential account of phenomenal concepts is due to Brian Loar (1997), and that deserves special attention. Phenomenal concepts, according to Loar, are recognitional; they apply in virtue of one's ability to recognize things of some relevant type. There are two worries about this. One is that, as with knowing 'wh', it's doubtful that one has the ability to recognize anything that one cannot describe in some way, even if that description is by itself often not precise enough to enable recognition of the relevant thing.

Loar considers recognizing a type of cactus. As he notes, one might typically have that ability without being able to describe the type of cactus sufficiently well to enable such recognition. Still, it's hard to imagine a case in which one could recognize a type of cactus and yet be unable to say anything descriptive that could help individuate that type.⁸

⁸ One might construe an ability to recognize as one's knowing how to recognize. And then knowing how to recognize something would require that one have some relevant descriptive knowledge about that type of thing. Still, that does not imply the controversial claim of Stanley and Williamson (2001) that knowing how is itself a kind of knowing that. See in the present context their discussion of the Mary case (pp. 443-4), as well as their discussion of Mary in terms of gaining a new skill (Stanley and Williamson, forthcoming).

It is in any case unclear that an ability to recognize is accurately seen in general either as a skill or as knowing how to recognize. One might be able to recognize something without there being any skill involved in doing so. And many abilities resist being described as cases of knowing how to do something; it's at best misleading to describe one's ability to look off to the left as one's knowing how to look off to the left.

So having the ability to recognize a case of one's seeing red should similarly require that one be able to describe, even if in some minimal way, what it is to be a case of seeing red. And if we understand phenomenal concepts as applying in virtue of a recognitional ability, applying the phenomenal concept of seeing red will require being able to describe, using ordinary, nonphenomenal concepts, a case of seeing red. Construing phenomenal concepts in terms of recognitional abilities does not avoid the demand to specify what it is for one to see red in ordinary, nonphenomenal, descriptive terms.

And there is a second difficulty. Mary's seeing red might very well enable her to recognize subsequent cases of seeing red, but there can be no guarantee. Seeing a sample of a particular type of cactus might enable one to recognize future cases of that type, but it might not. Similarly, Mary's knowing that she is in a particular token state cannot ensure that she would be able to recognize future instances of that type, or even that she has any grasp, recognitional or otherwise, of what type that particular token belongs to.

Seeing red is so familiar to us that it's tempting to think that it's unmistakable; if one has seen red, surely one will recognize any future case. But extrapolating from our situation to Mary's is unwarranted; her first experience of seeing red might be so overwhelming in its novelty as to result in its being unclear whether future chromatic experiences are of the same type. And if Mary were to see red without gaining the ability to recognize future cases, she would not have a phenomenal concept of seeing red construed in recognitional terms. I'll return to and expand on this point in §IV, along with other issues that arise in connection with the appeal to phenomenal concepts.

III. Ability and Acquaintance

It is unlikely that phenomenal concepts, however construed, will help provide an account of the conceptual content of the new factual knowledge that Mary is supposed to gain on first seeing something red. This is especially so if phenomenal concepts are understood in terms of a recognitional ability, since there might be nothing conceptual at all in the mere gaining of an ability.

This invites consideration of a proposal by Lawrence Nemirow (1980, 1990, 2007) and David Lewis (1983, 1999) that what is new to Mary is not factual knowledge at all, but rather just an ability to recognize, imagine, or remember the relevant type of experience.⁹ And there is on this account no new type of concept that results from gaining the relevant ability.

Similarly, if you can recognize somebody, it's also arguably at best misleading to describe that as your knowing how to recognize that person.

⁹ Nemirow relies principally on the ability to imagine (1990, 495).

This ability hypothesis seeks to accommodate the idea that Mary comes to have knowledge not covered in her books, but unlike the appeal to phenomenal concepts it denies that this new knowledge is factual. Since the new knowledge doesn't show there are nonphysical facts, it can't undermine physicalism.

But the considerations that caused difficulty for Loar's proposal apply equally to the ability hypothesis. For one thing, Mary might on first seeing red simply not gain any of the relevant abilities. One can have an experience of a new type but be unable to recall its nature or imagine or recognize another instance of its type.

But even if Mary did gain the relevant recognitional ability, she would also gain factual knowledge about the kind of state she comes to be able to recognize. This is clear from the foregoing discussion of recognitional concepts; it is difficult to believe that one could come to be able to recognize something and yet be unable to say anything descriptive, however general, about the thing one comes to be able to recognize. And such descriptive information would be factual.

One might suppose that this factual knowledge would be new to Mary, since it accompanies her new recognitional ability. But we cannot simply assume that; the recognitional ability might be new but the factual knowledge that new ability elicits might have been covered in her books. To assess whether the factual knowledge that accompanies Mary's new recognitional ability could have been in her books, we would have to get clear about the descriptive content of that factual knowledge.

Whatever factual knowledge accompanies a new recognitional ability that Mary gains will describe the things Mary has come to be able to recognize, namely, the conscious qualitative character of her experiences. Though the ability hypothesis denies that gaining a recognitional ability involves gaining factual knowledge, there is compelling reason to hold that it does. So we need an independent account of the nature of the qualitative character of experiences to determine what factual knowledge might accompany Mary's first seeing something red or gaining the ability to recognize such an experience. Only then can we assess whether such factual knowledge could have been in Mary's books.

There is another proposal that, like the ability hypothesis, seeks to accommodate Mary's gaining new knowledge but deny that the new knowledge is factual. On that proposal, the new knowledge is a type of acquaintance; "learning what an experience is like is identical to becoming acquainted with the experience" (Conee 1994, 140). Following Bertrand Russell's famed distinction between descriptive knowledge and knowledge by acquaintance (1912, ch. 12), Mary might then gain new knowledge by becoming newly acquainted with an experiential property, but not thereby gaining any new factual knowledge. In the same spirit, Paul Churchland urges that Mary's new knowledge might just be "knowledge by acquaintance" (1985, 24), and hence not factual.

But as with the proposal that the knowledge Mary gains is an ability to recognize, it's doubtful that one can become acquainted with something without having some descriptive knowledge relevant to that thing. Being acquainted with something in the ordinary case

involves perceiving it, and that in turn involve must some descriptive knowledge about it. Acquaintance with an experience is due in some way to consciousness, not perception. But there is no reason to think that the same doesn't hold there. Being acquainted with a conscious experience always enables one to say something informative about that experience, because being acquainted with it rests on some measure of descriptive information about it. Some descriptive knowledge is always built into acquaintance.¹⁰ What we need to know is whether that could have been in Mary's books.

One might suppose that consciousness results in acquaintance with our experiences in some direct way that precludes any descriptive knowledge. Indeed, some such picture very likely underlies the appeal to phenomenal concepts. I'll turn to that idea in the next section. But that picture aside, there is no reason to suppose that acquaintance with anything, even with conscious qualitative experiences, can occur without one's being able to describe what one is acquainted with, if not precisely and in detail, at least in general terms.

The ability and acquaintance hypotheses are proposals for seeing Mary as gaining new knowledge, but not descriptive knowledge and so not knowledge that could be covered in her books. But in both cases, the new nondescriptive knowledge would be accompanied by some relevant descriptive knowledge. Only when we know exactly what content that descriptive knowledge would have can we assess whether it could have been in Mary's books.

IV. What Consciousness Tells Us

The appeal to phenomenal concepts is intended to ensure that on first seeing something red Mary comes to have knowledge that is both new to her and factual. That the knowledge would be new is built into what phenomenal concepts are; one cannot have them without having a token of the relevant type of qualitative experience. The difficulty is in ensuring that the so-called phenomenal concepts are genuinely conceptual, since unless they are the new knowledge won't be factual.

Still, it may continue to seem inviting that one does get new factual knowledge of some sort when one has a qualitative experience of a new type, knowledge that is not merely one's acquiring a new ability or one's coming to be acquainted with an experience of that type. Consciousness, it may seem, tells one something new when one has a qualitative experience of a type new to one. And what consciousness tells us, it may seem, is factual; one comes to have factual information that one didn't have before.

On this picture, consciousness tells one, without mediation or help from any other source, what it's like for one to have each conscious experience, what the qualitative character is

¹⁰ This goes beyond Russell's mild concession that "it would be rash to assume that human beings ever, in fact, have acquaintance with things without at the same time knowing some truth about them" (1912, p. 46), and explains why what he says is so.

of any experience one has. It must be without help from any other source to ensure that the factual knowledge consciousness delivers will not be in Mary's books. Consciousness on its own gives one factual knowledge about the qualitative character of each experience.

But it's by no means obvious what factual information that could be. Consciousness plainly does often provide one with an experience of a type new to one. But that by itself does not show that consciousness also by itself provides one with factual knowledge about those new experiences. Simply providing a conscious experience of a type new to one will not sustain the knowledge argument.

And the new factual information one gains cannot simply be that this, referring to the new experience, is what it's like to see something red. The knowledge that this is what it's like to see something red refers, by hypothesis, to an experience that's new to one. But it does so demonstratively or by picking the experience out in relation to other current experiential states one can describe. Though the experience referred to is of a type new to one, the referring does not by itself constitute new informational content. The referring in one's knowledge that this is what it's like to see something red does not provide new conceptual content.

The factual knowledge that this is what it's like to see something red does not simply refer to an experience of a type new to one; it also predicates a property of that experience. But it's not obvious that that property that factual knowledge predicates of the experience referred to could not have been described in Mary's books. The factual knowledge says that the experience in question is a case of what it's like to see something red. Why couldn't that property be in Mary's books?

It's tempting to say it couldn't be in Mary's books because the property of being what it's like to see something red is ineffable. So that property is not subject to being described in any articulate way at all. Consciousness in effect shows us the property, so that we know from consciousness what that property is. But it simply is not the sort of property that could have been in Mary's books, because it is not the sort of property that lends itself to being described in any way at all. This is reflected in Ned Block's colorful appeal, in saying what qualitative mental states are, to Louis Armstrong's remark about jazz: "If you gotta ask, you ain't never gonna get to know" (Block 1978, 281).

This is why the appeal to phenomenal concepts is so inviting. If the knowledge that this is what it's like to see something red involves applying a phenomenal concept of the experience that knowledge refers to, then the first time one has an experience of the sort that phenomenal concept applies to there will be conceptual content that is new to one. So the factual knowledge will itself be new to one. The appeal to phenomenal concepts is a way of capturing the idea that though consciousness does tell us something about our qualitative experiences, what it tells us cannot be put in words; it does not allow for any articulate description.

This fits with the earlier discussion of a problem for phenomenal concepts. The difficulty was in establishing that phenomenal concepts so called are genuinely conceptual, that

they contribute to the conceptual content of the new factual knowledge Mary is supposed to gain. Absent any conceptual connections that phenomenal concepts would have with other relevant concepts, there is reason to doubt that phenomenal concepts are concepts at all. And the lack of such conceptual connections goes hand in hand with the apparent ineffability of what consciousness is supposed to tell us about our qualitative experiences. If the so-called phenomenal concepts lack ties to ordinary concepts, there will be no way to articulate the content of the alleged phenomenal concepts.

This, to reiterate, is because phenomenal concepts have as modes of presentations one's own qualitative experiences, the very experiences that the phenomenal concepts are posited to apply to in a distinctively first-person way. But it is difficult to see how such concepts could have robust conceptual ties with any other concepts, and without such robust ties phenomenal concepts cannot figure in any articulate or informative remarks even about one's own qualitative experiences. Hence the ineffability.

Those who posit phenomenal concepts may be tempted to urge that they operate only partly as ordinary nonphenomenal concepts do, and that the demand that they connect conceptually with other concepts is accordingly unreasonable. But the advocate of phenomenal concepts must then ensure that the nonstandard nature of phenomenal concepts not only allows them to qualify as a kind of concept at all, but also and more important that their nonstandard nature as concepts will permit them to do the job the advocates have in mind. It is far from clear that either demand can be met.

It might seem, however, that consciousness does somehow give us articulate information about our qualitative experiences after all. Consider an experience of seeing something red; doesn't consciousness tell us that this experience is more like an experience of seeing something orange than like an experience of seeing something blue? Doesn't consciousness tell us countless other comparative things about the various types of qualitative experience we have?

In general, comparing things requires being able to identify and individuate those things independently of the comparison. If one compares two objects in respect of their size or shape or color, one must be able to pick those things out independently of the comparison. So for consciousness to deliver comparative information about qualitative experiences, we must be able, relying exclusively on what consciousness tells us, to identify and individuate the experiences being compared independently of the comparisons. So such comparisons tell us nothing about the nature of the individual mental qualities that we wouldn't already have known from what consciousness had told us about each mental quality on its own. At least this is so if consciousness is to deliver comparative information without help from any other source. And if consciousness gets help from another source, whatever information that other source contributes might well be in Mary's books.

Consciousness can compare qualitative experiences only if it has some way to identify and individuate each experience on its own. But it's unclear how consciousness might do that. Consciousness by itself seems to have nothing informative to say about individual experiences; whatever consciousness might tell us about them is ineffable, incapable of

being described. So consciousness by itself cannot deliver comparative information about our qualitative experiences. Nothing that consciousness by itself could tell us about our qualitative experiences could appear in Mary's books. But that is not because consciousness gives us information of some special sort. It's because consciousness tells us nothing about our qualitative experiences that could figure in any factual knowledge about them.¹¹

V. Qualitative Experience

The idea that we can learn what the nature of qualitative experiences is solely from what consciousness tells us is basic to the knowledge argument and to the positing of phenomenal concepts. Phenomenal concepts are consciousness-based concepts, whose possession and application rest solely on how consciousness presents the relevant qualitative character. It's this feature of phenomenal concepts that seals them off from conceptual ties with any other concepts.

The idea that we can learn what the nature of qualitative experiences is solely from what consciousness tells us is often simply taken for granted in current the philosophical literature, as though it's so obvious that no argument or assessment is called for. But without some articulate account of what information consciousness provides us about qualitative experiences, it's hard to avoid the conclusion that consciousness, by itself, doesn't provide any information at all. The fallback claim that what consciousness tells us is ineffable doesn't help; if it's ineffable, it isn't information. It certainly isn't anything that could underwrite the idea that Mary on first seeing something red gains new factual knowledge. If it's factual, it can be expressed descriptively. One might contest that connection, but it is difficult to see what being factual could amount to otherwise.

Still, it may seem that there is no alternative to this consciousness-based picture, that there simply is no source of information about the nature of qualitative experience other than consciousness.¹² So it's important to assess whether there is an alternative to the exclusive reliance on consciousness that underlies the knowledge argument. Is there a way to tell about the nature of qualitative experience and about specific types of qualitative experience other than from the inside?

¹¹ Thomas Nagel's (1974) has famously urged that conscious experience requires to be treated in subjective terms, and that subjectivity drives out any objective treatment. This is a natural conclusion if a treatment is subjective only if it relies solely on consciousness. But one need not so construe subjectivity; one can construe it in a way that both does justice to the contrast between subjective and objective and also leaves open the possibility of an objective treatment of conscious experience (Rosenthal 1983).

¹² This conviction reflects the influence of Nagel (1974), and is nicely epitomized in Adam Pautz's evocative slogan, 'consciousness first' (2013, p. 195). It is also reflected in David Chalmers's claim that "there is nothing we know more directly than" mental qualities, that qualitative character "is the most vivid of phenomena; nothing is more real to us" (1996, 3).

The widespread reliance on first-person access may blind to the crucial tie that qualitative experiences have with perception, a tie that we can exploit to construct an alternative to the consciousness-first approach. Our having qualitative experiences is intimately tied to our ability to discriminate among stimuli of various sorts; indeed, it is in virtue of our having qualitative experiences of different types that we are able to discriminate perceptible stimuli at all. One discriminates stimuli with distinct shades of red by having different qualitative experiences of the relevant stimuli, and similarly with all perceptual discrimination.

Since qualitative experiences are the states that enable us to discriminate perceptible stimuli, we should be able to use the tie between qualitative experiences and perceptual discrimination to understand both the nature of qualitative experiences and how to characterize the various types of qualitative experience. Whenever two stimuli are discriminable by an individual, that individual will have qualitative experiences that differ themselves in a way that corresponds to the relative difference among the stimuli.

This correspondence of differences between stimuli and qualitative experiences can be made precise. We can begin with the most minimal case of perceptual discrimination, so-called just noticeable differences (JNDs), where an individual is able to discriminate two stimuli that would be indistinguishable were they physically any closer. Experimentally, display to an individual two stimuli that are physically identical; the individual will report that they're the same. Then adjust one of the stimuli to be slightly different from the other in its physical nature by successive minimal steps until the individual reports that now they're different. Those two stimuli are now JND.¹³

Having established JNDs among all stimuli accessible by a particular modality, we can then use these minimal discriminable differences among stimuli to construct a space that represents the distance in respect of JNDs between any two stimuli thus accessible. Since differences in qualitative experiences enable these discriminations, we have a space that

¹³ There are various methodological issues in this procedure that aren't relevant for present purposes. Experimental subjects tends to be conservative, so that they may continue to report no difference even when they can discriminate the stimuli; this can be controlled for by interspersing irrelevant displays. And noise in the perceptual channels may yield inconsistent reports, requiring statistical techniques to get a definite result. There are also issues about how to handle the well-known intransitivity of JNDs; see Goodman (1961, pp. 256-258). JNDs must be tested individual by individual, but the results can be averaged to get discriminability for a species or other group.

It is crucial for present purposes that what is just noticeably different are the perceptible stimuli and not, as some discussions in psychophysics assume, the subjective assessments of differences in mental quality between distinct experiences. That's because the present proposal is to explain qualitative experiences by appeal to their role in the perceptual discrimination of stimuli.

also represents the qualitative distance between any two experiences. So the space of JNDs among discriminable stimuli also provides us with an account of the mental qualities of experiences that are responsible for such discriminability (Rosenthal 1991, 1999, 2001, 2005, chs. 5-7, 2010).¹⁴

This quality-space theory of qualitative character and of its various types makes no appeal to what consciousness might tell one about qualitative experiences generally or about particular types. It relies exclusively on testing the ability to discriminate various stimuli and the tie between discriminative ability and types of qualitative experiences. It's plain that if on being presented with two stimuli the qualitative experiences that result are the same in type, one will be unable to perceive any difference between them. Only when qualitative experiences differ can we distinguish the stimuli that give rise to them.

The quality space of mental qualities derives from the quality space of discriminations that an individual can make among various stimulus properties. Constructing that quality space relies on testing what discriminations are possible, which won't always follow the physical nature of those stimulus properties. This is as it should be; types of qualitative character don't reflect the physical nature of stimuli, but the discriminative abilities an individual has in respect of those stimuli.¹⁵ Differences among of qualitative characters are what enable perceptual discriminations of stimuli.

The connection between experience and perceptual discrimination is fundamental to our conception of what a qualitative experience is. Quality-space theory simply capitalizes on that connection, and constructs a way to assess and calibrate the discriminable differences among stimuli with differences among mental qualities. And because the quality-space account of qualitative character makes no appeal to what consciousness might tell one, it provides a robust alternative to the consciousness-first approach, which at best delivers ineffable pronouncements devoid of articulable, factual information.

VI. Consciousness vs. Perceptual Role

Quality-space theory, which appeals to the role qualitative states have in perception, offers a clear, articulate account of what qualitative character is and how the various types of qualitative character are individuated. Exclusive reliance on consciousness to learn about qualitative character, by contrast, provides no descriptive knowledge, and so nothing that

¹⁴ For more on quality spaces generally, see Clark (1993) and Kuehni (2010).

This quality-space account of mental qualities can be extended to provide a way of individual the perceptual modalities themselves; two mental qualities belong to the same modality if there is a chain of JNDs that lead from one to the other (Young, Keller, and Rosenthal 2014, Rosenthal 2015).

¹⁵ So quality-space theory differs crucially from Churchland (2007), which relies on the physical nature of the relevant stimuli.

could sustain the idea that Mary gets new knowledge about what it's like for one to see something red simply by having a new conscious experience of seeing something red.

Still, the exclusive reliance on consciousness is so entrenched in current discussions of qualitative character that it's worth addressing objections that those who endorse that reliance may raise for quality-space theory, as well as considering several additional ways that quality-space theory has advantages over a consciousness-based approach.

One concern a consciousness-first theorist will immediately raise is the neutrality of quality-space theory about whether the mental qualities in question are conscious. Perceptual discrimination occurs without being conscious; indeed unconscious perception can yield very accurate discriminations; forced-choice guesses in subliminal perception of various types, for example, in masked priming (Marcel 1983, Cheesman and Merikle 1986, Breitmeyer and Öğmen 2006), are typically highly accurate. And change detection can occur in the absence of subjective awareness (Fernandez-Duque and Thornton 2000, Laloyaux 2003).¹⁶ If mental qualities are, as quality-space theory hypothesizes, the mental properties responsible for perceptual discrimination, then mental qualities occur in these cases without being conscious.

Unconscious qualitative character is anathema to a consciousness-based approach, on which we learn about qualitative character exclusively from consciousness. And many discussions simply identify qualitative character with what Block (1995) calls phenomenal consciousness, in effect building consciousness into qualitative character by terminological stipulation. But absent some substantive reason to insist that qualitative character cannot occur without being conscious, we should resist that stipulation.¹⁷

One might take the term 'mental quality' as simply implying consciousness; it is likely that the term 'quale' does carry that implication, as does Block's (1995) term, 'phenomenal consciousness'. But even if one so construed 'mental quality', one can always adopt another term that is neutral between being conscious and not. And construing the term as implying consciousness would in any case beg the substantive question of whether the mental properties we refer to in the conscious case as mental qualities can also occur without being conscious.

¹⁶ Indeed, unconscious discrimination is sometimes more accurate than conscious discrimination (Scott and Dienes 2010).

¹⁷ A traditional way to differentiate perception from nonperceptual cognition is that perception involves mental qualitative character, whereas cognition on its own does not (cf. Firestone and Scholl 2016, p. 1). But this relies on conscious perception, and so cannot work if qualitative character cannot occur without being conscious but perception can. This is doubtless responsible for the rejection by many today of that traditional line between perception and cognition, and the search for an alternative that does not appeal to qualitative character at all (e.g., Block 2016). Accommodation of qualitative character that need not be conscious would obviate the need for that search.

An advocate of that approach might insist that the insistence that qualitative character cannot occur without being conscious is not mere stipulation, but reflects compelling pretheoretic intuition. But pretheoretic intuition would be unable to access unconscious qualitative states if they did occur; intuition can at best tell only about those mental phenomena that are conscious. So the appeal here to pretheoretic intuition at best merely channels the consciousness-based approach.

In any case, so-called intuitions often serve simply as cover for undefended theoretical assumptions, for getting others to accept theoretical claims without having to defend them.¹⁸ And relying on intuition here seems an especially vivid case of that, since whether qualitative states can fail to be conscious is a theoretical, not a commonsense issue. We cannot settle the question by appeal to intuition, but must appeal to relevant theoretical considerations.

And many considerations, both theoretical and commonsense, favor quality-space theory and its appeal to the role qualitative character plays in perceptual discrimination. Suppose you see an object of a particular shade of red and I ask you to tell me about your conscious visual experience. One way you might reply would be to tell me your experience is more like that of seeing a red tomato than that of seeing a fire engine or the like. These objects needn't be present when one sees the target object; they are objects whose characteristic colors we readily remember and so can compare with the shade under consideration.

Such a reply is along the lines of quality-space theory; you're describing the relative location your visual experience has in a space that is fixed, in turn, by a range of discriminable red stimuli, stimuli that typically exhibit characteristic shades of red. Such a reply can't sustain an appeal to consciousness. You're not telling me something you learn from consciousness, but about how your experience compares with other experiences you might have of objects known to us both.

There is another possible reply to such a question. If there are several red objects visible in the current scene, you might instead tell me how your experience of the target red object compares to your experiences of those others. That is again in keeping with quality-space theory; you're locating the target experience within a space of experiences fixed in turn by stimuli we can independently pick out. This second reply does not rely just on something consciousness tells us about the nature of the target experience. In commonsense contexts we describe our qualitative experiences comparatively, by how they compare to other experiences that we can pick out by appeal to independently identifiable stimuli.

¹⁸ As what we might call theory pumps, inverting Daniel Dennett's inviting idea of an intuition pump (1980, 429; 1991, 282, 397).

The consciousness-based approach has notable historical precedents, e.g., among British empiricists and among logical empiricists. But the view was advanced by those writers as a theoretical claim, to be defended by overall theoretical considerations, not by appeal to allegedly pretheoretic intuition.

It might seem that terms like 'red,' 'green', and 'blue', which readily anchor types of visual experience, are not comparative after all. And if they are not, perhaps their use rests on some intrinsic qualitative character as revealed by consciousness. But that impression is misleading. If we describe the qualitative character of seeing red to somebody who is red-green colorblind, we would explain those ostensibly noncomparative terms comparatively. The experience of seeing something red, for example, is typically more like an experience of seeing yellow than like one of seeing blue. Even ostensibly noncomparative terms, such as 'red' and 'green', get cashed out comparatively when we can't rely on shared experience. Our pretheoretic, commonsense ways of describing mental qualities are comparative, relying on relevant similarities and differences.

Indeed, the way color input is processed neutrally itself fixes the color mental qualities comparatively, and not one by one, even for colors that figure dominantly in our classification of things, such as red, green, and blue. Information about levels of excitation of rods and cones in the retina is passed onto cells that process that information comparatively, comparing the strength of red stimulation with that of green stimulation, blue with yellow, and white with black. This so-called opponent processing results in the neural signals that result in mental qualities of color. Mental qualities of color are produced by comparisons that serve in effect to locate those mental qualities within a space of possible qualities fixed by opponent-processing comparisons. Mental qualities of color are not produced atomistically, using a distinct process for each type of mental quality.

Our commonsense practice of describing qualitative experiences by appeal to comparisons with experiences that result from a currently accessible or remembered range of stimuli undermines the idea that we know about qualitative experiences by what consciousness tells us about them. We know about them comparatively, by what types of stimulus they reflect. And if what we know about our qualitative experiences is not from consciousness, we no longer have any reason to reject the idea that qualitative character can occur without being conscious.

A striking aspect of our qualitative experiences is that we can identify and distinguish them in a far more fine-grained way when occur simultaneously than when they occur in succession. If one has several experiences of shades of red that are very close to one another, they may seem indistinguishable if they occur in succession but clearly different if they occur together (Halsey and Chapanis 1951, Pérez-Carpinell *et al* 1998). Indeed, we are aware of such qualitative experiences as indistinguishable when they occur in succession and as plainly distinct when they occur simultaneously. Consciousness relies on comparisons to individuate our qualitative experiences. This holds not only for color experiences, but also for the qualitative character distinctive of other modalities (e.g., Burns and Ward 1982). What consciousness tells us about our qualitative experiences is itself comparative.

Because consciousness presents shades in respect of more fine-grained differences when simultaneous than when in succession, some have claimed that the effect is due simply to limitations of memory, limitations that prevent one from comparing successive experiences

(e.g., Raffman 1995). But whatever explains the effect, being presented with a red stimulus of a particular shade results in a different subjective experience when presented together with other stimuli with closely related shades as against being presented with no such stimuli. The availability of comparison with closely related stimuli changes the subjective character of the experiences that result. Consciousness presents the qualitative character of those experiences comparatively. Consciousness does not reveal intrinsic properties of experiences, which would be independent of any such comparisons.

The comparative way we ordinarily think about qualitative character, which reflects the way consciousness presents those experiences, gives reason to question the currently fashionable idea that the mental qualities of our experiences cannot diverge from the way consciousness presents them.¹⁹ Being presented with a particular color stimulus results in a different subjective experience depending on whether that stimulus is accompanied by others closely related to it. Imagine seeing a particular sample of red by itself, and then the same sample accompanied by one only very slightly different. Consciousness presents the two cases differently; the comparison available in seeing the two slightly different reds results in fine-grained subjective experience of each.

But the stimulus is the same in either case. So apart from subjective appearances, we have not reason to think each stimulus results in the same mental quality, regardless of whether the stimulus is accompanied by another. Consciousness accordingly presents us with subjective mental appearances that diverge from the underlying mental reality.²⁰ Subjective awareness can be wrong about the mental qualities that figure in perceptual discrimination (Rosenthal 2011; Peels 2016).

A particularly nice demonstration is Raffman's (2011) finding that when adjacent hues differ by less than just-noticeable differences, though subjects of course judge two to be indistinguishable, they adjust a third hue to match only one of the hues in a systematic way. Visual processing is more fine-grained when it relies on manual matching than on subjective awareness. In a different context, Scott and Dienes (2010) have shown that unconscious perceptual processing can be more accurate than conscious processing.

The denial that experiential states can subjectively appear differently from the way they actually are, which is seldom if ever argued, is simply a reflection of the adoption of a first-order theory of consciousness, since on such a theory there is no second factor that can diverge from the first-order experience.

It is possible that commonsense and ordinary usage stress first-person access in ways that seem to support the view that mental qualities cannot occur without being conscious. But the emphasis in ordinary conversational interactions on the first person likely reflects

¹⁹ Thus Nagel: "The idea of moving from appearance to reality seems to make no sense" in connection with conscious experience (1974, 444; cf 448). Cf. Kripke on pain (1980, 151-2).

²⁰ For additional considerations in support of that distinction, see Rosenthal (2018).

nothing more than the social impropriety of challenging first-person pronouncements about a person's mental states, which downplays our third-person access to others' states. This has no bearing on whether the distinguishing mental properties of conscious qualitative states can also occur without being conscious, and cannot outweigh the theoretical considerations adduced above.

Since mental qualities are the mental properties responsible for perceptual discriminations, the space of mental qualities mirrors that of discriminable stimulus properties. So we can take each mental quality to represent the stimulus property whose relative location in its quality space corresponds to the relative location of that mental quality in its quality space.

Representationalism is the view that the only qualitative character one is ever aware of in perceiving are properties one perceives something to have (Harman 1990, Byrne 2001). The only qualitative character that occurs in consciously seeing a green object, for example, is the green color one sees that object to have. There is no qualitative character in perceiving except the properties that perceiving represents perceived objects as having.

Quality-space theory defines a representational role for mental qualities. But quality-space theory is not a type of representationalism. Mental qualities are the distinguishing mental properties of perceptual states, and are independent of any representational role. We taxonomize and individuate qualitative states and explain their nature not by appeal to their representational role,²¹ but their role in perceptual discrimination. Indeed, the only sound reason for taking mental qualities to represent stimulus properties is the role they play in perceptual discrimination. Discriminative role forges the tie mental qualities have to stimulus properties, and mental qualities represent only because of that tie.

Because representationalism sees conscious qualitative character as solely a matter of the properties we perceive things to have, it rejects the compelling folk view that what it's like for one in conscious perceiving is a matter of the type of psychological state one is in. Quality-space theory does justice to that folk view by construing qualitative character as a matter of the perceptual mental states themselves.

Representationalism is appealing in part because it sidesteps mental qualities altogether, and so can avoid quandaries that mental qualities allegedly lead to. But quality-space theory can retain mental qualities and still dispel those quandaries by allowing mental qualities to occur without being conscious. Representationalism also trades on the inviting idea that the properties one perceives are in some way relevant to the nature of perceptual states. Quality-space theory has a richer explanation of how the properties we perceive figure in the nature of perceptual states, again without jettisoning mental qualities; mental qualities are individuated by discriminability relations among the stimulus properties.

²¹ As Jacob Berger (2018) in effect urges we do, in developing and arguing for an ingenious and novel form of representationalism. My thanks to Berger for suggesting the title for this paper.

Since quality-space theory relies on perceptual role for an account of qualitative character, it's neutral as regards whether states with qualitative character are conscious. So we need additional resources to explain why some qualitative states are conscious and others not. It's useful to begin with the commonsense observation that if we have good reason to conclude that an individual is in some qualitative state but that individual is wholly unaware of being in it, that state is not a conscious state. And since being wholly unaware of a mental state that one is actually in is sufficient for that state not to be conscious, a necessary condition for a state to be conscious is that one be aware of being in it (Rosenthal 2005; cf. Dienes 2004 for relevant empirical tests).

Being aware of a state requires being in some state in virtue of which one is aware of that target state.²² But that higher-order state need not itself be conscious. The qualitative state does not inherit its being conscious from the higher-order state's being conscious; rather one's being aware of oneself as being in the qualitative state itself constitutes that state's being conscious. Since the way we're aware of our conscious states seems subjectively to be unmediated, it must not seem subjectively that this higher-order awareness relies on inference or observation. I have argued elsewhere that the way one is aware of a mental state that is conscious is by having a thought to the effect that one is in that state (e.g., Rosenthal 2005). That fits well with our ability to report being in states that are conscious; those verbal reports express the relevant higher-order thoughts. But for present purposes, we can stick in more generic terms simply with an appeal to some suitable higher-order awareness.²³

Still, what's crucial for present purposes is that this account of what it is for mental states to be conscious explains how subjective awareness presents states as having different mental qualities. When a qualitative state is conscious, one is aware of oneself as being in a state with some particular type of mental quality. But what it is for a mental qualities to be of a particular type is for it to have a particular relative location in a quality space.

²² One might posit that the target state makes one aware of itself (e.g., Kriegel 2009), but that makes for unnecessary complications in accommodating qualitative states that aren't conscious. Also, there is evidence that the onset of qualitative states precedes their becoming conscious (Libet 2004, ch. 2). And there are other disadvantages of a theory that builds the higher-order awareness into the state (Phillips 2007; Rosenthal 2004, §5).

We cannot construe one's being conscious of one's conscious states as simply a cognate accusative, as with one's smiling one's smiles (Sosa 2003, p. 276), since that would preclude cases in which one fails to be conscious of one's mental states.

²³ So a state is conscious if one is aware of it and not otherwise. So being conscious is dichotomous in that way, not graded (contra Overgaard *et al* 2006). Nonetheless, the higher-order awareness in virtue of which a state is conscious can itself be more or less vivid or compelling. And that suggests that states that are conscious can be conscious to different degrees.

So one's being aware of being in a state with a particular type of mental quality simply is being aware of oneself as being in a state with a mental quality that occupies a particular location in the relevant quality space. Consciousness presents mental qualities comparatively, in respect of their relative quality-space location. This explains why consciousness presents states resulting from closely related stimuli as indistinguishable if they occur successively but different if they occur simultaneously; consciousness presents the states in respect of relative location in a quality space. It is unlikely that any other account can explain this effect as well. Consciousness does not present mental qualities in respect of some intrinsic property of the qualitative state.

It may seem that an account of the nature of conscious qualitative character in terms of relative location in a quality space, supplemented by a higher-order awareness of being in a state with such a relative location, is overly intellectual and in that way does not comport with our pretheoretic intuitions about qualitative experiences. It may be thought that since those intuitions proceed at a commonsense level, a highly theoretical account cannot conform to them.

But that's a mistake. The desire for an account that does nothing but conform to our pretheoretic intuitions stems from assuming that consciousness is our only source of knowledge about qualitative experiences. If we know about qualitative experiences in other ways, as I've argued we must, we have to expect that those other ways will involve a measure of theorizing.

The correct constraint is that the theorizing conforms to our commonsense way of describing qualitative experiences, even if it also goes well beyond it. And quality-space theory and a higher-order theory of consciousness satisfy that demand. We describe qualitative experiences in comparative terms, and we describe what it's like to have a particular experience by being aware of ourselves as having particular experiences and comparing them with others we can independently identify by appeal to characteristic or currently present stimuli.

VII. Mary's Knowledge

The knowledge argument requires that Mary, on first seeing something red, come to have knowledge about what it's like that is both factual and new. And the foregoing discussion poses difficulties for Mary's having knowledge that is not only new but also factual.

Mary does come to know that she is having an experience unlike any she has had before. That is new factual knowledge, but it's not the knowledge needed for Jackson's argument. And it is plainly comparative knowledge, comparing her new experience with experiences she has had before. By itself, it is knowledge only that her new experience differs from previous ones; it is not even knowledge of how it compares with those she's had before.

Suppose, following most discussions, that one adopts a consciousness-based approach to how one learns about one's qualitative experiences; consciousness tells one directly all

one can know about the nature of those experiences. It's compelling, on such an approach, that Mary does get new knowledge. The trouble is that there seems to be no way to cast that new knowledge as factual. We cannot express descriptively what we learn on a consciousness-based approach. The new knowledge is ineffable.

This is why many have pursued the possibility that the new knowledge is either a matter of acquaintance or a matter of gaining a recognitional or imaginative capacity. Even if one opts for phenomenal concepts to capture Mary's new knowledge, the most influential version of that hypothesis (Loar 1997) itself appeals to recognitional abilities. And it's likely that this version tacitly underlies other attempts to formulate Mary's new knowledge by appeal to phenomenal concepts.

A consciousness-based approach encourages us to think that Mary's knowledge on first seeing something red is new but causes difficulty in construing that knowledge as factual. But we have seen that a perceptual-role approach, such as quality-space theory, has advantages over a consciousness-based approach. How does Mary's knowledge fare on that kind of view?

Mary sees something red for the first time, and she says, "This is what it's like to see something red." That statement expresses knowledge about her new experience; she is characterizing that experience in respect of its conscious qualitative character. Quality-space theory provides a straightforward factual account of exactly what Mary's conscious qualitative character consists in. It is qualitative character that occupies a particular relative location in the quality space of mental color qualities.

And Mary is also aware of being in a state with qualitative character specified in respect of that relative location, a higher-order awareness of the first-order qualitative state. On quality-space theory, Mary's knowledge of what it's like to have an experience of seeing red has clear descriptive content. When she says or thinks, "This is what it's like to see something red," the descriptive content of her statement or thought is that her experience compares to other color experiences in such-and-such a way.²⁴

But Mary is a special case. Before first seeing something red, her visual experiences had all been grayscale. So the range of color experiences she has available to compare her new experience with is vastly more limited than the experiences with which we could compare such a new experience. The quality space in terms of which one could describe the comparisons available to Mary would be far smaller than the quality space that normal individuals operate with.

So what it would be like for Mary to have her new experience would not be what it's like for us to see red. Indeed, since the quality space in terms of which Mary can compare

²⁴ Churchland suggests a similar view when he writes, rather in passing, that Mary's knowledge about sensations and their properties might "be a matter of being able to make certain sensory discriminations, or something along these lines" (1985, 23; emphasis Churchland's).

conscious experiences is vastly smaller than ours, it is unclear that we can form any accurate picture of what it would be like for her. Mary would not learn what it's like to see red, in general; she would not even learn what it's like simply for her to see red, since that too would change with exposure to stimuli of various versions. If she does learn what it's like at all, I would at best be what it's like for her to see red having previously been confined to grayscale stimuli.

One might urge that Mary does, after all, have a normal quality space. Mental qualities are fixed by the discriminations among stimuli that an individual is capable of making. And if presented with a range of chromatic hues, perhaps Mary could discriminate among them much as we do. Construing Mary's discriminative ability to include new stimuli she could be presented with would reflect her potential discriminative ability. And it would yield a quality space much like ours, not limited to those she has already seen.

On this more liberal construal of discriminative ability, Mary would have mental qualities corresponding to all the discriminations she could make if presented with the relevant stimuli. This seems theoretically extravagant. Mary's visual apparatus has the potential to generate those mental qualities, but in advance of being presented with the relevant stimuli she doesn't yet have them.²⁵ The quality space that figures in what it's like for Mary is a matter of experiences she has had, not those she could have in presented with the relevant stimuli.

But even if Mary did have all those latent mental qualities awaiting suitable stimuli, there would still be nothing it's like for her to be in states with the unactivated mental qualities. Mary would never be aware of herself as being in such states until she is actually in such states. The latent mental qualities are simply mental qualities her perceptual apparatus provides her with the potential to have; they are not actual mental qualities. So Mary would be unable subjectively to compare her new experience with those she would have if presented with novel stimuli. So far as conscious subjectivity is concerned, Mary is restricted to comparing her new experience with others she has actually had.

Since the subjective color comparisons available to Mary are far fewer than ours, the descriptive content of her thought or remark, "This is what it's like to see red," will be different from the content of anything we might say or think. Mary's content will be roughly that the new experience is very different from anything she has ever seen, along with some comparisons with various grayscale shades.²⁶

By contrast, when we see something red and say or think, "This is what it's like to see red," our descriptive content involves locating our experience among a dramatically greater

²⁵ Indeed, the operative areas of visual cortex will likely have been recruited for other functions. That may be reversible, but likely not immediately. See, e.g., Kauffmann *et al* 2002.

²⁶ This is not an artifact of Mary's prior limitation to grayscale stimuli. The point would hold equally if she has previously had grayscale and green stimuli available.

range of colors. Mary's first experience of seeing red will be subjectively very different from ours and likely very different from ours even independent of subjective considerations. It may be difficult if possible at all for us to imagine subjectively what her experience would be like, though there is no difficulty at all in describing it.

It's useful to consider an argument of Lewis's against a particular type of appeal to resemblance relations. "A literalist," he writes, "might see the phrase 'know what it's like' and take that to mean: 'know what it resembles'. Then he might ask: ... Why can't you just be told which experiences resemble one another? You needn't have had the experiences—all you need ... is some way of referring to them." But as Lewis notes, "'know what it's like' does not mean 'know what it resembles'." He concludes: "If you are taught that experience A resembles B and C closely, D less, E not at all, that will help you know what A is like—if you know already what B and C and D and E are like. Otherwise, it helps you not at all" (1999, 365-6).

When Mary first sees something red, she has a new experience. And we can colloquially say that in having that experience she knows what it's like to have it. The phrase, 'know what it's like', can refer in this minimalist way to one's simply having the relevant experience. So understood, knowing what it's like is not a kind of knowing at all. So on that minimalist construal, it's overly literalist to see the phrase, 'know what it's like', as involving any kind of knowing at all.

But knowing what it's like can also, more generously, mean actually having knowledge about what kind of experience one has. And on that construal, resemblance relations do help. The only factual knowledge²⁷ available to Mary or to anybody about the kind of experience Mary comes to have is knowledge about the resemblance relations that define the relative location of that experience in Mary's color quality space. This is not because of some literalist reading of 'like' in the phrase, 'what it's like'. It's simply because there is no other way to give a genuinely informative description of any type of experience. There are no other terms in which one can express factual knowledge about the kind of qualitative experience anybody has.

Jackson's argument rests on Mary's coming to have new factual knowledge, knowledge that could not have been in Mary's books. A consciousness-based approach permits thinking that Mary does get new knowledge, though it's hard to see how it could be factual. Quality-space theory, by contrast, casts Mary as having knowledge that is plainly factual. The issue is whether that knowledge is new, whether it could have been in Mary's books.

What it's like for one to see something red is a matter of one's being aware of oneself as having an experience of a particular type. On quality-space theory, one is aware of oneself as having an experience with a particular relative location in a suitable quality

²⁷ Lewis, as noted in section III, urges that we understand this knowledge not as factual, but in terms of a recognitional ability. But as argued there, recognitional abilities bring with them factual knowledge, knowledge, expressible in sentential form.

space. Mary's awareness of that relative location will be limited to color experiences she's already had. Still, there's no difficulty in constructing the quality space that her subjective comparisons rely on; it's all the grayscale shades plus the new shade of red.

And that quality space could readily be in Mary's books, along with color quality spaces for normal individuals and those with various deficits, such as one or another type of color blindness. And since what it's like for Mary to see something red for the first time is a matter of the qualitative comparisons she could be subjectively aware of, what it's like for her could itself readily be in Mary's books, as well as what it's like for the rest of us.

Jackson holds that Mary's books, not being historical accounts, would not contain knowledge of what it's like for Mary to see red for the first time, but insists that she would, on first seeing something red, gain new factual of what it's like for people in general to see something red. But there's no difficulty in Mary's books containing both types of factual knowledge. And reading Jackson's article would doubtless have motivated the authors of Mary's books to be sure to include both types of factual knowledge. Neither type of factual knowledge would be new to Mary.

I've argued that quality-space theory is preferable on many counts to a consciousness-based approach. But the core difficulty for Jackson's argument is a dilemma that doesn't require choosing between the two approaches. On a consciousness-based approach, the knowledge Mary has on first seeing something red is new, but not factual. On quality-space theory, whatever knowledge Mary has is factual, but not new; it would be in Mary's books.

There is of course something new to Mary when she first sees something red; she has her first conscious experience of seeing something red. But what is new is not factual knowledge. An account that relies on perceptual-role, such as quality-space theory, provides that Mary can have knowledge that is factual about the kind of experience she comes to have, but that factual knowledge will be in her books and so not new. A consciousness-based approach may encourage the idea that Mary comes to have knowledge of some sort that would not be in her books, but on such an approach it is unclear what such knowledge could be. There is in any case no reason to see Mary as coming to have any knowledge that is both new and factual.

The idea that Mary gains new factual knowledge likely derives from running together these two incompatible ways of understanding the nature of qualitative experience. On one of way of understanding the nature of qualitative experience, Mary has knowledge that's factual about her new experience. On the other, that knowledge is new to her. But it's only the experience that's new to Mary; she gains no factual knowledge about the experience that is also new to her.

Mary has mastered all the factual knowledge in her books. So she has factual knowledge about what it's like for somebody whose previous color experiences are all achromatic to see something red for the first time. Would her first seeing something red elicit that factual knowledge?

It might well. Compare having factual information about what somebody looks like. Seeing that person for the first time might elicit the factual knowledge one had. But it also might not. One's factual knowledge about what the person looks like might not be sufficiently detailed for the experience to elicit that knowledge. Or that knowledge might simply not come to mind.

Similarly, Mary's books can contain factual knowledge about what it's like for somebody whose previous color experiences are all grayscale to see something with a particular shade of red. But that experience might not, for whatever reason, elicit the factual knowledge from Mary's books. It is natural to expect that it might well not. Mary might simply be at a loss as to just what kind of new experience she is having.

But the knowledge that concerns Jackson is knowledge of what it's like for people in general to see something red. That's knowledge about the type of experience it is. If that knowledge is factual, it will be in Mary's books. Mary comes to have no new factual knowledge on first seeing something red, and hence none that causes any problem for physicalism.²⁸

²⁸ My thanks to Sam Coleman for numerous comments, which led to many significant clarifications.

REFERENCES

- Alter, Torin, and Sven Walter (2007), eds., Phenomenal Concepts and Phenomenal Knowledge: New Essays on Consciousness and Physicalism, Oxford: Oxford University Press.
- Berger, Jacob (2018), "A Defense of Holistic Representationalism," Mind & Language, 33, 2 (April): 161-176. doi: 10.1111/mila.12163
- Block, Ned (1978), "Troubles with Functionalism," in Minnesota Studies in the Philosophy of Science, IX, ed. C. Wade Savage, Minneapolis: University of Minnesota Press, 1978, pp. 261-325.
- Block, Ned (1995), "On a Confusion about a Function of Consciousness," The Behavioral and Brain Sciences, 18, 2 (June): 227–287.
- Block, Ned (2016), "Tweaking the Concepts of Perception and Cognition," Behavioral and Brain Sciences, 39: 21-2. doi:10.1017/S0140525X15002733
- Breitmeyer, Bruno G., and Haluk Ögmen (2006), Visual Masking: Time Slices through Conscious and Unconscious Vision, 2nd edn., Oxford: Oxford University Press.
- Burns, Edward M., and W. Dixon Ward (1982), "Intervals, Scales, and Tuning," in The Psychology of Music, ed. Diana Deutsch, New York: Academic Press, pp. 241-269.
- Byrne, Alex (2001), "Intentionalism Defended," The Philosophical Review, 110, 2 (April): 199-240.
- Chalmers, David J. (1996), The Conscious Mind: In Search of a Fundamental Theory, New York: Oxford University Press.
- Cheesman, Jim, and Philip M. Merikle (1986), "Distinguishing Conscious from Unconscious Perceptual Processes," Canadian Journal of Psychology, 40, 4 (December): 343-367.
- Churchland, Paul M. (1985), "Reduction, Qualia, and the Direct Introspection of Brain States," The Journal of Philosophy, 82, 1 (January): 8-28.
- Churchland, Paul M. (2007), "On the Reality (and Diversity) of Objective Colors: How Color-Qualia Space Is a Map of Reflectance-Profile Space," Philosophy of Science, 74, 2 (April): 119–149.
- Conee, Earl (1994), "Phenomenal Knowledge," Australasian Journal of Philosophy, 72, 2 (June): 136-150.
- Clark, Austen (1993), Sensory Qualities, Oxford: Clarendon Press, 1993.

- Dennett, Daniel C. (1980), "The Milk of Human Intentionality," Behavioral and Brain Sciences 3, 3 (September): 428-30.
- Dennett, Daniel C. (1991), Consciousness Explained, Boston: Little, Brown and Company.
- Dienes, Zoltán (2004), "Assumptions of Subjective Measures of Unconscious Mental States: Higher-Order Thoughts and Bias," Journal of Consciousness Studies, 11, 9: 25-45.
- Farkas, Katalin (2015), "Knowing-wh Does Not Reduce to Know-that," American Philosophical Quarterly 53, 2 (April): 111-124.
- Fernandez-Duque, Diego, and Ian M. Thornton (2000), "Change Detection without Awareness: Do Explicit Reports Underestimate the Representation of Change in the Visual System?," Visual Cognition, 7, 1-2-3 (January–March): 324–344.
- Firestone, Chaz, and Brian J. Scholl (2016), "Cognition Does Not Affect Perception: Evaluating the Evidence for 'Top-Down' Effects," Behavioral and Brain Sciences, 39: 1-19. doi:10.1017/S0140525X15000965
- George, B. R. (2013), "Knowing-'wh', Mention-some Readings, and Non-Reducibility," Thought 2, 2 (June): 166-177.
- Goodman, Nelson (1951), The Structure of Appearance, Cambridge, Massachusetts: Harvard University Press.
- Halsey, Rita M., and Alphonse Chapanis (1951), "On the Number of Absolutely Identifiable Spectral Hues," Journal of the Optical Society of America, 41, 12 (December): 1057-1058.
- Harman, Gilbert (1990), "The Intrinsic Quality of Experience," Philosophical Perspectives, IV: 31-52
- Higginbotham, James (1996), "The Semantics of Questions," in The Handbook of Contemporary Semantic Theory, New York: Blackwell, pp. 361–83.
- Jackson, Frank (1982), "Epiphenomenal Qualia," The Philosophical Quarterly, 32, 127 (April): 127-136.
- Jackson, Frank (1986), "What Mary Didn't Know," The Journal of Philosophy, 83, 5 (May): 291-295.
- Kauffmann, Thomas, Hugo Théorot, and Alvaro Pascual-Leone (2002), "Braille Character Discrimination in Blindfolded Human Subjects," NeuroReport 13, 5 (April): 571-574.
- Kriegel, Uriah (2009), Subjective Consciousness: A Self-Representational Theory, Oxford: Oxford University Press.

- Kripke, Saul A. (1980), Naming and Necessity, Cambridge, Massachusetts: Harvard University Press.
- Kuehni, Rolf G. (2010), "Color Spaces and Color Order Systems", in Color Ontology and Color Science, ed. Jonathan Cohen and Mohan Matthen, Cambridge, Massachusetts: MIT Press, 3–36.
- Laloyaux, Cédric, Arnaud Destrebecqz., and Axel Cleeremans (2003), "Implicit Change Identification: A Replication of Fernandez-Duque and Thornton (2003)," Journal of Experimental Psychology: Human Perception and Performance, 32, 6 (December): 1366–1379.
- Lavie, Nilli (1995), "Perceptual load as a Necessary Condition for Selective Attention," Journal of Experimental Psychology: Human Perception and Performance 21, 3 (June): 451-68.
- Lavie, Nilli (2010), "Attention, Distraction, and Cognitive Control under Load," Current Directions in Psychological Science 19, 3 (June): 143-148.
- Lavie, Nilli, Diane M. Beck, and Nikos Konstantinou (2014), "Blinded by the Load: Attention, Awareness, and the Role of Perceptual Load," Philosophical Transactions of the Royal Society B: Biological Sciences, 369, 1641 (May 5): 20130205, 1-10.
- Lewis, David (1983), Postscript to "Mad Pain and Martian Pain," in Lewis, Philosophical Papers, vol. 1, New York: Oxford University Press, pp. 130-132.
- Lewis, David (1999), "What Experience Teaches, in Lewis, Papers in Metaphysics and Epistemology, Cambridge: Cambridge University Press, pp. 262-290; originally in Proceedings of the Russellian Society, Sydney: University of Sydney, 13 (1988): 29-57
- Libet, Benjamin (2004), Mind Time: The Temporal Factor in Consciousness, Cambridge, Massachusetts: Harvard University Press.
- Loar, Brian (1997), "Phenomenal States," in The Nature of Consciousness: Philosophical Debates, ed. Ned Block, Owen Flanagan, and Güven Güzeldere, Cambridge, Massachusetts: MIT Press/Bradford Books, pp. 597–616.
- Ludlow, Peter, Yujin Nagasawa, and Daniel Stoljar (2004), eds., There's Something about Mary: Essays on Frank Jackson's Knowledge Argument, Cambridge, Massachusetts: MIT Press/A Bradford Book.
- Lycan, William G. (1996), Consciousness and Experience, Cambridge, Massachusetts: MIT Press/Bradford Books.
- Marcel, Anthony J. (1983), "Conscious and Unconscious Perception: An Approach to the Relations between Phenomenal Experience and Perceptual Processes," Cognitive Psychology 15, 2 (April): 238-300.

- Nagel, Thomas (1974), "What Is It Like to Be a Bat?", The Philosophical Review 83, 4 (October): 435-50.
- Nemirow, Laurence (1980), review of Thomas Nagel, Mortal Questions, The Philosophical Review, 89, 3 (July): 473-477.
- Nemirow, Laurence (1990), "Physicalism and the Cognitive Role of Acquaintance," in Mind and Cognition: A Reader, ed. William G. Lycan, Oxford: Basil Blackwell. pp. 490-499.
- Nemirow, Laurence (2007), "So This Is What It's Like: A Defense of the Ability Hypothesis," in Alter and Walter (2007), pp. 32-51.
- Nida-Rümelin, Martine (2009), "Qualia: The Knowledge Argument," The Stanford Encyclopedia of Philosophy (Summer 2015 Edition), ed. Edward N. Zalta, URL = <http://plato.stanford.edu/archives/sum2015/entries/qualia-knowledge/>
- Overgaard, Morten, Julian Rote, Kim Mouridsen, and Thomas Zoëga Ramsøy (2006), "Is Conscious Perception Gradual or Dichotomous? A Comparison of Report Methodologies during a Visual Task," Consciousness and Cognition 15, 4 (December): 700-708.
- Pautz, Adam (2013), "Does Phenomenology Ground Mental Content?", in Phenomenal Intentionality: New Essays, ed. Uriah Kriegel, Oxford: Oxford University Press, pp. 194-234.
- Peels, Rik (2016), "The Empirical Case against Introspection," Philosophical Studies, 173, 9 (September): 2461-2485.
- Pérez-Carpinell, Joaquín, Rosa Baldoví, M. Dolores de Fez, and José Castro (1998), "Color Memory Matching: Time Effect and Other Factors," Color Research and Application, 23, 4 (August): 234-247.
- Phillips, Ben (2014), "Indirect Representation and the Self-Representational Theory of Consciousness," Philosophical Studies, 167, 2 (January): 273-290.
- Pylyshyn, Zenon W. (2007), Things and Places: How the Mind Connects with the World, Cambridge, Massachusetts: The MIT Press/A Bradford Book.
- Raffman, Diana (1995), "On the Persistence of Phenomenology," in Conscious Experience, ed. Thomas Metzinger, Exeter, UK: Imprint Academic, pp. 293-308.
- Raffman, Diana (2011), "Vagueness and Observationity," in Vagueness: A Guide, ed. Giuseppina Ronzitti, Dordrecht: Springer, pp. 107-121.
- Robinson, Howard (1982), Matter and Sense: A Critique of Contemporary Materialism, Cambridge: Cambridge University Press.

- Rosenthal, David (1983), "Reductionism and Knowledge," in How Many Questions?, ed. Leigh S. Cauman, Isaac Levi, Charles Parsons, and Robert Schwartz, Indianapolis: Hackett Publishing Co., pp. 276-300.
- Rosenthal, David (1991), "The Independence of Consciousness and Sensory Quality," in Philosophical Issues: Consciousness, 1: 15-36.
- Rosenthal, David (1999), "The Colors and Shapes of Visual Experiences," in Consciousness and Intentionality: Models and Modalities of Attribution, ed. Denis Fisette, Dordrecht: Kluwer Academic Publishers, pp. 95-118.
- Rosenthal, David (2001), "Color, Mental Location, and the Visual Field," Consciousness and Cognition, 9, 4 (April): 85-93.
- Rosenthal, David (2004), "Varieties of Higher-Order Theory," in Higher-Order Theories of Consciousness, ed. Rocco J. Gennaro, Amsterdam: John Benjamins Publishing Company, pp. 19-44.
- Rosenthal, David (2005), Consciousness and Mind, Oxford: Clarendon Press.
- Rosenthal, David (2010), "How to Think about Mental Qualities," Philosophical Issues: Philosophy of Mind, 20 (October): 368-393.
- Rosenthal, David (2011), "Exaggerated Reports: Reply to Block," Analysis 71, 3 (July): 431-437,
- Rosenthal, David (2015), "Quality Spaces and Sensory Modalities," in The Nature of Phenomenal Qualities: Sense, Perception, and Consciousness, ed. Paul Coates and Sam Coleman, Oxford: Oxford University Press, pp. 33-65.
- Rosenthal, David (2018), "Seeming to Seem," in The Philosophy of Daniel Dennett, ed. Bryce Huebner, New York: Oxford University Press, pp. 133-164.
- Russell, Bertrand (1912), The Problems of Philosophy, London: Oxford University Press.
- Russell, Bertrand (1927), An Outline of Philosophy, London: George Allen & Unwin, Ltd.
- Scott, Ryan B., and Zoltán Dienes (2010), "Knowledge Applied to New Domains: The Unconscious Succeeds Where the Conscious Fails," Consciousness and Cognition, 19, 1 (March): 391-398.
- Sosa, Ernest (2003), "Privileged Access," in Consciousness: New Philosophical Perspectives, ed. Quentin Smith and Aleksandar Jokic, Oxford: Clarendon Press, pp. 273-94.
- Stanley, Jason, and Timothy Williamson (2001), "Knowing How," The Journal of Philosophy 98, 8 (August): 411-444.

Stanley, Jason, and Timothy Williamson (forthcoming), "Skill," Noûs, advance online version at <http://onlinelibrary.wiley.com/wol1/doi/10.1111/nous.12144/full>

Stoljar, Daniel (2016), "The Semantics of 'What it's like' and the Nature of Consciousness," Mind 125, 500 (October): 1161-1198.

Vendler, Zeno (1972), Res Cogitans: An Essay in Rational Psychology, Ithaca, New York: Cornell University Press.

Young, Benjamin D., Andreas Keller, and David Rosenthal (2014), "Quality Space Theory in Olfaction," Frontiers in Psychology (January), <http://journal.frontiersin.org/article/10.3389/fpsyg.2014.00001/full>.