

SECONDARY QUALITIES: WHERE CONSCIOUSNESS AND INTENTIONALITY MEET

1

When giving an introductory lecture on the mind-body problem, I have standardly presented it as consisting of two basic questions: how can a material being possess intentional states, and how can there be something it is like to be in a material state? That the two issues of qualitative, or phenomenal character, and intentionality are distinct had been a consensus view for many years. However, recently that consensus has broken down under pressure from various quarters.

Perhaps the main source of pressure is from representationalism. For the representationalist, the phenomenal character of a conscious sensory state is constituted by the representational, or intentional content of that state (Byrne 2001, Dretske 1995, Harman 1990, Rey 1998, and Tye 2000). So what it's like to see a red surface, on a ripe tomato say, is just to represent that surface as red.

One way, then, of joining the problems of phenomenal character and intentionality is to reduce the former to the latter. This is what the representationalist proposes, and of course, as a way of solving the mind-body problem, it has clear merit. If you can reduce two hard problems to one, even if that one is pretty hard in its own right, you've surely made progress. What's more, for a time at least, philosophers were pretty confident that they had a handle on intentionality, so prospects looked bright for a general resolution of the mind-body problem. Also, aside from the obvious gain in contracting the problem space, the idea, following Brentano, that intentionality is the mark of the mental, gave the view that phenomenal character really just is a kind of intentional content an air of great plausibility. That mental phenomena should be unified in this way

makes for a much more satisfying theory than we get if the two problems—phenomenal character and intentionality—are treated separately.

For reasons I won't rehearse here (see Levine 2003), I don't think one can provide a reductionist account of phenomenal character in terms of intentional content. However, I do believe that the two phenomena are intimately connected. For one thing, though I don't accept the reductionist consequences drawn by representationalists from this consideration, I do find what they call the argument from "transparency" somewhat convincing. The point is that, when we reflect on what it is to experience phenomenal character, as in the case of my seeing a ripe red tomato on the kitchen counter, it certainly seems right that I'm being presented with a way the world is. What it's like for me does seem essentially a matter of how things appear to me, and how things appear to me certainly seems to involve intentional content.

If one doesn't buy the representationalist-reductionist program, but one still acknowledges the intimate link between phenomenal character and intentional content, then what emerges is a view on which, instead of looking to intentionality to illuminate phenomenal consciousness, we find phenomenal consciousness seriously complicating our view of intentionality. This is what I propose here; I'm afraid it's just the opposite of progress, assuming one measures progress by how closely one approaches a solution to the mind-body problem.

It seems to me that there are two routes into this complicating picture, and they both lead to a view of intentionality—or at least a variety of intentionality—that is fundamentally different from the standard one. What they lead to is an "awareness/appearance" (AA) model, as opposed to what I'll call a "representational" model.¹ On a representational model, intentionality is realized in the mind by means of the mind's embodiment—in us, the brain—containing symbolic tokens of some sort, and their intentional content derives from their causal relation to what they are about together with their structural relations to each other. The Language of Thought hypothesis is the obvious paradigm example, but what I have in mind is meant to cover pictorial and connectionist models as well. The essential elements are these: the mind is in an intentional state by virtue of tokening a representation, what the intentional state is directed on is a function of the content of the representation in question, a token's having the content it has can be functionally analyzed in non-intentional (proba-

bly causal) terms, and, perhaps most importantly, the relation itself does not involve consciousness in an essential way.

On the AA model, the intentionality of the relation is not a function of the intentional content of a token representation, intentionality is not functionally analyzable, and consciousness is built into the relation. The subject—not a symbol—stands in a relation of awareness, conscious awareness, to the object, while the object stands in an appearance relation to the subject. The intentionality is built directly into the relation; it's not a property of a symbol as in the representation model. I'll have more to say about the nature of this relation later, but this should give the flavor of what I have in mind.

One of the two routes to the AA model starts from consideration of phenomenal concepts. Phenomenal concepts are those we employ when considering our phenomenal experiences from the first-person point of view. Many philosophers who want to make phenomenal experience safe for materialism appeal to the special nature of these concepts in order to explain our strikingly resilient intuitive resistance to materialist accounts of the phenomenal (Balog 1999, Levin 2007, Loar 1997, Papineau 2002, and Perry 2001). It seems to me that only if we adopt something like the acquaintance model of first-person access to experience can this appeal work, but then it becomes just as much a problem for materialism as was phenomenal experience in the first place. I won't try to make this case here as I've discussed it at length elsewhere (Levine 2007).

What I do want to consider is the second route, which begins with consideration of secondary qualities. The problem of how to locate secondary qualities like color and sound in the world exercised early modern philosophers of course, but it has become a popular topic nowadays as well. Mostly the discussion has concerned the nature of color, but of course many of the same issues arise with the other secondary qualities. I think the AA model suggests itself after consideration of the problems that beset the representational model. The idea is that secondary qualities are essentially ways of appearing to a conscious subject, and thus not mind-independent properties of external objects. Elaboration of this claim will have to wait until presentation of the motivation behind it.

In Sections 2 and 3 I will present two lines of reasoning that motivate adoption of the AA model of secondary qualities. The first comes out of the recent literature on the status of color and its relation to color expe-

rience, and the second brings together the controversy over the possibility of unmet pains and the old saw about the tree in the forest. I hope to show that a consistent line of reasoning emerges from these different strands that gives us reason to take seriously what is admittedly a difficult model of conscious intentionality to swallow. Finally, in Section 4 I explore what modifications in the AA model are necessitated by the existence of hallucinatory experiences.

2

Let me begin my discussion of the first motivation by listing what seem to me the "data" that a reasonable theory of color and color experience begins with (though not necessarily ends up with). First, color seems to be a property of objects. We certainly apply color terms to objects, and the reasonable assumption is that the truth-maker for an assertion such as "the apple on the table is red" is that the apple in question instantiates a certain property, whatever property it is that 'red' designates. Reinforcing this sense of objectivity for color is the phenomenon of color constancy. We see objects as possessing the same color despite obvious changes in appearance under different lighting conditions. We have a sense that some conditions (say daylight) better reveal an object's color to us than others (say fluorescent light). Color constancy thus adds to the sense that the true color is a property of the object, not just how it appears to us.²

Second, both color itself and our concept of it have an especially intimate relation to our visual experience of color. With regard to the latter, it is common for people to say that congenitally blind people, though they can certainly learn to use color vocabulary, and presumably make meaningful assertions with it, still don't really have a full concept of color. The idea is that you don't really know what red is unless you have seen it. Why should that be, if color is an objective property of physical objects? True, vision is our primary means of epistemic access to color, but so long as there are other ways to know about color—and there are—why should inability to utilize the primary way most people have count against one's "fully possessing" the concept?³ The fact that we do find it plausible to deny blind and (totally) color-blind people full possession of the concept of color suggests a special relation between color and color experience.

As for color itself, the intimate relation to color experience is evident if we ask about its causal role. After all, what difference does an object's

color make to anything other than how it appears to a subject with color vision? Undoubtedly, due to co-evolution, certain plants propagate by virtue of their color, since bees and other insects, responding to the color, provide the mechanism by which seeds are spread. But the fact remains that it's the effect on a creature's visual system that is the principal and direct means by which color serves this function.⁴

Third, the relation between an object's color and how it looks to normal human perceivers is grounded in a very complex neurological mechanism that manifests a high degree of relativity of response to several parameters, such as ambient light, contrast with surrounding surfaces, state of adaptation of the relevant neural circuits, and the like. It is this fact in particular that has bedeviled various attempts to identify a well-behaved physical correlate of color in the world. Surface spectral reflectance, a disposition to reflect certain percentages of various wavelengths of light, is the best candidate, but even this runs into trouble with relativity. If color is really out there in physical objects, it's a very complex property of them indeed.

Fourth, color, if it is anything objective, is a physically constituted (or realized) property. Physicalism about the mind is of course a controversial doctrine, and whether color experience in particular can be explained in physical (which includes functional) terms is still a matter of intense debate. But few partisans of dualism about qualia want to make the same move with respect to color. If color is real and mind-independent (the latter a big "if" in my book), it's physically realized.

Finally, five, color experience, like all phenomenal experience, presents the world to us as being a certain way. In visual experience, in particular, the world *looks* a certain way. In other words, phenomenal experience has an essentially intentional character. What it's like to see something is usually captured by saying how that thing looks, which seems to be a way of saying what properties it looks to have. This so-called "transparency" argument, as remarked above, has been very influential in fueling representational theories of phenomenal character, and, as I mentioned, I myself find it quite compelling, though I don't accept all the consequences that many philosophers draw from it.

So now two interconnected questions face us: first, what is color? and second, what property is it we are presented with in color experience? With regard to the first question, the initial divide is between color realism (or objectivism) and color irrealism (or subjectivism). Realism has (roughly) three versions: physicalism, dispositionalism, and primitivism. Physicalists

identify color with a physical property of the stimulus, such as surface spectral reflectance (SSR). Dispositionalists identify color with what can be called an "appearance property," a disposition to appear to a subject in a certain way under certain conditions. For primitivists, color just is what it is, an objective property of the stimulus but not reducible to any other property. Irrationalism is the view that color isn't really out there in the world, but rather is a matter of our occupying, as Hardin (1988) puts it, "chromatic states."

Now primitivism is pretty clearly ruled out by number four, the claim that color is physically realized. But why accept this constraint? The main reason is that we know, whatever color is and whatever color experience is, that the causal mechanism by which the mind apprehends color is physical, so if color is a property of the object that isn't realized physically, color vision becomes a miracle.

Irrationalism, though I think it has a lot going for it actually, suffers from *prima facie* implausibility. First, one has to swallow the consequence that nothing out there in the world is really colored. So the things we say, like "that apple is red," aren't literally true. But in addition, this error theory must extend to perception itself, due to datum five above. That is, if transparency is right, then color experience, like all phenomenal experience, presents the world as being a certain way. In color experience, it clearly presents the world as being colored. But if nothing is colored, we are forced to say that in this respect experience systematically misleads. While adopting an error theory for color judgments and experience isn't perhaps the end of the world, many would agree it's to be avoided if possible.⁵

This leaves the remaining two forms of realism, physicalism and dispositionalism. I referred to dispositionalism above as the view that colors are appearance properties, because the natural way to understand the disposition in question is as the disposition to appear a certain way to a perceiver. One consideration that favors the appearance-property view over physicalism, at least *prima facie*, is datum two above. If the only real effect of color in the world is on the way objects appear to perceivers, then it seems to make sense to identify color with this property of disposing the object to appear a certain way. However, this consideration isn't by any means decisive, since one might take account of this datum by treating the characterization of color in terms of a disposition to appear a certain way as the primary reference fixer embodied in our concept of color, but not

actually what the concept refers to. This is how Hilbert (1987), who describes his position as "anthropocentric realism," treats color.

In fact, there is one clear advantage physicalism has over the appearance-property account in characterizing the disposition to appear a certain way. The problem is how to describe the manifestation of the disposition, the way things appear. One might have thought one could just appeal to the qualitative state itself that is the effect of stimulation from the relevant object. So to be red, say, is to be disposed to cause a reddish sensation in a perceiver. But if we adopt transparency, that won't work, since having a reddish sensation involves perceptually representing the object as being a certain way. Now, if the way in question is describable independently of any disposition to appear a certain way, as would be the case if color is identical to an SSR, then there's no problem. But if color is itself just a disposition to appear a certain way, and that way is to appear colored, it seems as if the property isn't properly grounded.

Despite this problem, to which I will return shortly, there are good reasons for thinking that whatever color itself really is in the end, one has to posit something like an appearance property to capture the content of color experience.⁶ Here's why. Suppose visual color perception only represented an objective feature of the object, such as an SSR. In that case, when two subjects view the same surface but have different experiences, or when one subject views the same surface under different conditions and sees it differently, one would have to say that only one among the conflicting experiences was veridical. That is, one would have to say that if one treated phenomenal content as intentional content. So long as one can't appeal to merely intrinsic differences in the experiences of the two subjects (or the one subject at different times), so that the difference in the nature of the two experiences must be cashed out in terms of how the experience represents what is being seen, one can't have incompatible phenomenal characters both veridically representing the same situation. One of them has to be wrong.

The problem is that there doesn't seem to be a principled basis for calling any of them wrong, at least not one that doesn't apply to the others, so the phenomenal differences can't be differences in the objective properties being represented. But since they have to be characterized as a difference in representational content, there must be something else, besides

objective property (if objective property is there at all in the representational content), in the representational content of the experience.

So whether or not objective color is represented in experience, it seems we have to admit appearance properties as well into the content (so long as we hold on to transparency as well). An appearance property overcomes the problem of incompatible yet equally veridical experiences as follows. Suppose Alice and Bill are both looking at a green color chip, C1. To Alice it looks to be pure, unique green, with no hint of blue or yellow in it, but to Bill it has a slightly yellowish cast to it, though still clearly green. Now there are color chips, C2 and C3, such that C2 looks to Alice the way C1 looks to Bill, and C3 looks to Bill the way C1 looks to Alice. Both Alice and Bill distinguish among C1–C3 with respect to their color appearances. So if objective color were all that that their experiences represented, one of them would have to be wrong.

However, instead we say that the differences in experience are not differences in the objective color represented—either we say the experiences don't represent objective color at all, or that they represent the very same objective colors. Rather, Alice and Bill are experientially representing different appearance properties that C1–C3 have. Alice is perceiving C1 to have appearance property A1 while Bill is perceiving C1 to have appearance property B1. A1 and B1 are of course different properties, so this gives their experiences different intentional content, which is what we want (if we are to hold on to transparency). But unlike the case of objective color, where an object cannot simultaneously possess two distinct SSR's, when it comes to appearance properties there is no problem in saying that C1 possesses (or instantiates) both A1 and B1. Alice is representing it as having A1, while Bill is representing it as having B1. They differ in what their experiences attribute to the object, but not in a way that makes their experiences conflict. Both can be correct.

Let's take stock with respect to our initial data. SSR's are instantiated in distal objects in a mind-independent way, and they can serve either as the objective colors themselves, if we go physicalist, or as the categorical grounds of the colors, if we go dispositionalist. This takes care of datum one, that colors seem to be objectively attributable to objects. SSR's are obviously physically realized, so datum four is handled as well. Data two, three, and five—that there is an especially intimate relation between color and color experience, that color experience is determined

by very complex factors in addition to SSR, and that phenomenal character is intentional (transparency)—are handled by having phenomenal character identified with the representation of appearance properties. So are we done?

Not quite. What about the grounding problem mentioned above? Isn't it circular to define phenomenal character in terms of the representation of color appearance and then define color appearance as a disposition to cause phenomenal character? Perhaps not. Shoemaker (2001/2002) makes the argument that one can interdefine the two as a "package deal," much the way belief and desire are interdefined by functionalists. He puts it this way:

... we might try to define a particular state/property pair by saying that the state of being appeared blue to and the property of appearing blue are the unique state S and the unique property P, such that S represents P and something has P just in case it is producing ... a state of type S. . . . (page 467)

On Tuesdays, Thursdays, and Fridays I tend to think this will work; the other days, not so much. Given that today is Friday, I'm not going to quibble with his "package deal" method. So maybe circularity, strictly speaking, can be avoided. Still, there is something quite odd about thinking of appearance this way, isn't there?

I think the problem can be brought out as follows. If we look back at the way the discussion about color and color experience has developed, as briefly reviewed above, one sees a tension pulling theorists in opposing directions. On the one hand, color is presented to us as out there in the world. When we see red tomatoes it sure seems like it's the tomato that's red. What's more, as representationalists (at least of the externalist variety) have so well noted, it's by putting the qualitative character of experience out there in the world that you overcome a major obstacle to accepting materialism. On the other hand, given the actual facts about color and its intimate relation to color experience, we are constantly driven back inside. So the appearance-property account seems to resolve the tension, allotting both inside and outside their due.

Yet, what happened to the blue on this account? (Or, should I say, echoing the old Pepsodent commercial, on this account you'll "wonder where the yellow went"?) The package-deal definition can maybe overcome circularity, but of course it can't pin down what makes one of these

S/P pairs a *bluish* appearance. To get a determinate qualitative character, or quality of an object, one must pin it down at least on one end. Standard externalist representationalism does this by pinning it down externally; determinate quality is identified with SSR (say), and qualitative character is then identified with representing that. But for all the reasons above, this doesn't seem to work. So we need to pin it down on appearance, and then treat external quality as that which brings this about. This means, in the end, that determinate quality—what makes it the case that Alice (from our example above) is seeing it precisely the way she is while Bill is seeing it precisely the way he is—must be pinned down by features (whether physical or functional) internal to the observer.

Shoemaker himself is of course aware of this. He is satisfied that each side is given its due so long as we make phenomenal character representational while pinning the specific quality on the nature of one's internal state. One problem, somewhat independent of this discussion, is that I don't see how his quasi-functional method for identifying a specific type of phenomenal character can work (Levine 1988).

But even leaving that aside, I think it doesn't really resolve the tension in the end. On this view of bluishness, one can't genuinely capture the dual nature of phenomenal character, how it really simultaneously characterizes one's state of awareness and the object of awareness. The problem is that the standard model requires a strict cleavage between the representation and what it represents, which is then reflected in the cleavage between awareness and what it is awareness of. But the tension that fuels the debate arises just because this cleavage needs somehow to be overcome. To see even more starkly the need to overcome the cleavage between subject and object in the consciousness of secondary qualities, I turn now to the second route to the AA model mentioned above.

3

Often the first thing one is asked after telling someone that one is a philosopher—unless they just want to tell you how much they hated their one philosophy class in college—is the old one about the tree falling in the forest with no one around to hear it: does it make a sound? I used to politely chuckle and say nothing, but if pressed I would say that of course

it makes a sound, it's just that no one hears it. What's the problem? Unless one endorses something like Berkeleyan idealism, this seems like the obvious response, and I presume it is the consensus opinion on the question among professional philosophers in the Anglo-American tradition.

Still, I've never felt totally comfortable with that standard reply. Yes, of course the tree's falling causes a series of airwaves to propagate, and hearing is a matter of detecting these airwaves, so what it is one would detect, and thereby hear, if one were there, is precisely what is there even if one isn't there to detect it. But yet there does seem to be a gulf—a “gap” if you will—between the idea of there being airwaves that cause oscillations in the inner ear and the idea of there being a sound. Sound, as opposed to the mere propagation of airwaves, seems to be essentially a matter of how something, well, “sounds” to a conscious, sentient subject. No subject, no sound, it would seem, at least in one sense. But if one is hearing the tree fall, then it seems that what it is one is hearing must be independent of its being heard, which takes us right back to the original, standard answer to the puzzle. So what is the right way to characterize the situation?

Let's consider how Locke treated the question. One way to characterize Locke's position on the status of secondary qualities is that they are dispositions in external objects to cause certain sensations in us, where the individuating conditions of those sensations are independent of their representational content. So when Locke emphasizes that secondary qualities aren't really in the objects that appear to have them, what he's denying is not that the objects really possess the relevant dispositions, but rather that the manifestations of the disposition—what we experience when we see color and hear sound—is in them. The reddishness that I directly confront in experience is just a property of my mental state—or, in his terms, an “idea” in my mind.

What I find especially interesting about Locke's position is one of the arguments he employs to defend it. In response to someone who would insist that the quality is really in the object, not the mind, Locke compares color and sound to pain. He asks whether one would be inclined to attribute the pain experienced as a consequence of touching a red hot poker, say, to the poker itself, and answers that of course no one would be so inclined. But then, he wonders, what's so different about the non-painful warmth one feels when one touches the poker after it has sufficiently cooled? Why are

we inclined to attribute the warmth to the poker but the pain—which, after all, is just much more intense heat—to the mind?

Now, as an argument against a contemporary transparency advocate it isn't very convincing. For one thing, many representationalists about qualia argue that pains are in objects in just the way that colors are, it's just that the object in question is one's own body. Pain is the representation of the state of a part of one's body. The burning pain sensation one feels when grabbing a red hot poker is the detection of the burning of the skin in one's hand. The difference between the sensation of the warmth and the sensation of pain may just be a difference in the objects—the poker in the former case and the hand in the latter—to which the represented qualities are attributed.

Be that as it may, I'm not especially concerned with how convincing an argument this is, but with the suggestive analogy Locke draws between pain on the one hand and color and sound on the other. I think this analogy helps with the original puzzle about the tree in the forest. So, what do we say about the possibility of unfelt pains? That is, do we think it's possible for there to be a pain of which we are entirely unaware, not part of conscious experience at all? Unlike the case of unheard sound and unseen color, unfelt pain has been a very controversial issue. I want to suggest that Locke was right to compare secondary qualities to pain, and that unheard sounds and unseen colors should be controversial in the same way that unfelt pain is. To make the case, let's consider the problem of unfelt pain first.

Many philosophers who want to strictly distinguish between phenomenal qualities and our awareness of them—higher-order theorists (Rosenthal 1997, Lycan 1996) being the prime, but not only, examples—insist that it makes perfect sense to talk about pains that one doesn't feel. Feeling a pain, on this view, is a matter of being aware, having conscious access to, a quality which is itself ontologically independent of the awareness of it. Just as Locke wants to compare color to pain to make his point, these theorists want to compare pain to color to make their contrary point.

Such a view goes along with the position concerning pain mentioned above, that feeling pain is a matter of representing a certain quality, presumably damage of some sort, in a part of one's body. The thing felt, the pain, is the damage, and certainly one can have this property instantiated in one's body without being aware of it, and therefore, without "feeling"

it. Of course this makes perfect sense as far as it goes, but the question is whether merely having this damage-property instantiated in one's body is sufficient for having a pain.

Now one way of opposing the possibility of unfelt pains is to treat the question as basically a semantic one. One might claim, plausibly to my mind, that the way we use the term 'pain' makes it inappropriate to apply it in those cases where we totally fail to "feel" it, when there is no conscious awareness of it at all. David Lewis famously said that pain is essentially a feeling, and one way to interpret this claim is to treat it as a conceptual matter. It's just part of our concept of pain that one attributes it only when one is aware of it.

Insofar as this is just a semantic issue, there really is no substantive argument between the two sides regarding the metaphysical facts. The relevant damage properties are instantiable independently of any awareness of them, and there are states of mind that constitute awareness, or representation of these properties. Neither side disputes these facts. The only question is whether we use the term 'pain' when only the damage property is instantiated. Semantic questions are no doubt real questions, but when questions in philosophy of mind come down to them, when it's just a question of when it's appropriate to apply the term, I, for one, get very bored.

However, I don't believe this really is a semantic issue. When people intuitively resist the idea that there are unfelt pains, or when they agree with Lewis's statement that pain is essentially a feeling, I think they have something in mind far more substantive than when it's appropriate to apply the term 'pain'. Rather, the idea is that when you contemplate that very quality of, say, a toothache, the one that fills your awareness, it does not seem at all like the kind of thing that can be instantiated in one's tooth (or the nerve, or whatever) without one's being aware of it. That is, it is in its very nature a *feeling*, a moment of consciousness, something "for a subject." Put another way, properties like pain aren't ever merely instantiated; they are always, by their very nature, experienced. They are, if you will, modes of experience. If you take away the experiential aspect, the pain is gone too, though something very closely connected to the pain remains, namely the damage property that gives rise to the pain.

The idea, then, is that pain is just the kind of property that can only be instantiated as a way of being conscious, a determinate "what" in "what it's like" at a particular time. But that doesn't mean it isn't also an

appearance, a way one's body is being presented. Transparency theorists are right to insist that even bodily sensations have an intentional aspect; they present particular parts of our bodies as hurting, tickling, itching, etc. Where we differ is that standard representationalists take transparency to show that the quality—the pain or the itch—is independently in the object and merely represented in the mind. I want to say instead that pain is a way of being consciously aware of one's body, so that awareness is essential—and not merely in a definitional way—to the quality.

As I've pointed out above, both Locke and transparency theorists, though they disagree sharply about where to locate secondary qualities in the world, agree that one should treat colors and sounds the same way one treats bodily sensations like pains and itches. I concur. It seems to me that colors are, in the first instance—by which I mean that very quality immediately present in consciousness—a way something looks, where that it is being observed is essential to its instantiation. Sounds, too, are ways things, well, sound, to a conscious subject. In this respect the right answer to the tree-in-the-forest puzzle is that without anyone to hear it there isn't strictly a sound, but there is the categorical ground of the disposition to sound a certain way were someone there to hear it. Spectral reflectances can also be identified with colors in this derived way: that is, that they are the categorical grounds of the disposition to be seen a certain way (in particular lighting conditions, etc.) were a conscious subject of the right sort to see it.

Now one traditional way to make secondary qualities instantiable only in the context of experience is to attribute them to sense data. This would indeed explain why, in the relevant sense, ripe tomatoes aren't red when observers don't perceive them, since redness in this sense only applies to sense data, and without perception there are no sense data. However, though the metaphysics of sense data are anathema to most representationalists, this doctrine doesn't really diverge from the standard representational model of intentionality.⁷ It is still the case that we have the mind representing, or apprehending, a matter of fact distinct from the apprehending of it.

Instead, on the AA model I propose as an alternative, we should see awareness as an indissoluble relation between the conscious subject and the object appearing to the subject, with the quality as a property of the relation. That is, rather than analyzing the tomato's appearing red to me in terms of its causing a mental representation of its redness, we see the red-

ness as a modification of how it appears to me. Some objects stand in an appearing-redly relation to me, others in an appearing-bluely relation to me, and some parts of my body in an appearing-painful (or toothachy) relation to me. The idea is that the relation itself is intentional, not deriving its intentionality from a feature of a token representation in the mind.

Of course this idea has a medieval flavor about it, and seems to represent backsliding, rather than progress, on how to understand intentionality. (Remember, I did warn you up front.) But in fundamental ways it does seem to do justice to the way consciousness and intentionality are inextricably combined in the perception of secondary qualities. In Section 2 it emerged that neither pinning the determinate color quality on a feature of the object or the subject answered to both transparency and relativity. Well, pinning it on the character of the relation, and making that relation intrinsically intentional, seems to resolve that tension.

In this section we've seen a similar tension arise over the question whether secondary qualities can be instantiated independently of their being consciously apprehended. On the one hand, if awareness is "of" something, then that something should be independent of the awareness. On the other hand, unfelt pains, and unseen colors too, seemed not to really make sense. But if they are modifications of the awareness relation, then this unity in duality, as it were, is accounted for.

So what are the costs of adopting the AA model? The biggest one is figuring out how it can square with materialism. The move which seems to make intentionality tractable for materialism is precisely what is definitive of the standard representational model: the intentionality of a mental state/relation is reduced to the intentionality of a mental representation, an object that has its intentional content in virtue of its standing in a non-intentionally characterizable relation to its referent. But of course phenomenal consciousness already puts pressure on materialism, and if, like me, one is unhappy with extant materialist accounts of phenomenal consciousness, adding another problem for materialism isn't going to look like a major obstacle.

But does it even make sense? How can a relation like awareness/appearance be both qualitative and intentional in a primitive way? I admit to having qualms about this, but then I wonder: why not? True, if we operate under the constraint that intentionality cannot be basic in nature, then the reduction to a feature of a mental representation, and from there to some

causal relation, seems clearly the way to go. Further, we do have the model of language to make the idea of a symbol meaning something much more familiar. But I submit that if we don't operate under the materialist constraint, it isn't clear that there is any particular metaphysical location—whether in a symbol or a relation between subject and object—that is inherently more reasonable to locate primitive intentionality.

At any rate, I am content to explore the possibility of making such a model work for now, given what I consider the inadequacy of other accounts of the experience of secondary qualities and the consciousness-intentionality relation. However, there is one problem for the view that I would like to consider, and this is the topic of the next, and concluding section.

4

On the AA model, appearing-redly is an inherently conscious and intentional relation holding between the subject of experience and the object she is aware of. But what about hallucinations? If I'm looking at a red tomato, the appearing relation holds between the tomato and myself. This is all fine, as far as it goes, until we ask: What if there is no tomato there? What is the other relatum of the appearing relation in that case?

Disjunctivists in the theory of perception (e.g., Martin 2004) claim that when veridically perceiving an external object the object perceived is partly constitutive of one's perceptual state. This isn't quite the same view as the one being defended here, but it suffers from the same *prima facie* problem with hallucination. The problem for them is that if I'm visually hallucinating a tomato I seem to be in a state of the same mental type, yet there is no tomato there to be partly constitutive of that state, so how could it be partly constitutive even when it is there? Their answer is to deny the premise of the objection, that veridical perceptions and their internally indiscernible hallucinatory correlates are really states of the same "fundamental" mental type. Thus the tomato is indeed partly constitutive of the veridical-perceptual state, and this claim isn't undermined by the possibility of first-person indiscernible hallucinations of tomatoes.

There are two reasons I don't want to take this escape route. First, I find disjunctivism wildly implausible. I am not going to defend this claim here, but just note that it provides me with good reason for seeking a solution elsewhere. But second, it's not clear it would help for my purposes

anyway. I am trying to understand what redness is as we experience it. Insofar as hallucinations involve reddish experiences, some account has to be given of what a hallucinatory experience's being reddish consists in. Now, with my anti-disjunctivist impulses, I'd like the account to unify with the account I give of redness in cases of veridical experiences. So something must change in the account developed so far.

The dilemma I face is this. On the one hand, I want to treat appearance as a genuine relation, and to locate phenomenal character as a property of that relation. On the other hand, relations, when they hold, metaphysically entail the existence of relata, so the idea that the appearing relation holds when there's nothing to fill in for the second relatum seems not to make any sense. I am not confident that there is a solution to this problem, but in what follows I want to at least speculate about what shape a solution might take. I will end by drawing out some further consequences of the view for the question of the connection between consciousness and intentionality.

What is tempting in response to this problem is to appeal to merely intentional objects. After all, Brentano's original idea was that consciousness—the mental in general—is definable by its constituting a relation that can hold between the subject and things that don't exist. In the end I think this must be the way to go, but it isn't easy to see just how to make this work. In particular, what is supposed to be the metaphysical status of a merely intentional object?

One option is to think of merely intentional objects as occupants of other possible worlds. So when one is veridically perceiving a red tomato one is standing in the appearance relation to the red tomato that is actually out there, but when hallucinating a red tomato one is standing in the appearance relation to a merely possible red tomato. But there are three serious objections to this view. First, one has to posit merely possible objects. Of course modal realists like Lewis don't mind that, but it seems a good idea not to be forced into buying into that view of modality. Second, how does one come to stand in a relation of this sort to a merely possible object? Certainly not through any causal connection to it. Finally, even if there is a way to make sense of one's standing in this relation to a merely possible object, is this likely to be the very same relation that holds when veridically perceiving? It's hard to see how, and our task is to find a unified account of appearing that covers both veridical perception and hallucination.

So I don't think the appeal to merely possible objects is going to work. The most ontologically innocent way to characterize a merely intentional object is as a representation of that object. When I think about Santa Claus, say, one can explain what is meant by calling Santa Claus an intentional object in terms of my entertaining a very real mental representation of Santa Claus. To say that Santa Claus is the object of my thought is then just to say that I'm tokening a representation of Santa Claus. This is ontologically innocent in that the Santa Claus representation really exists, as just mentioned, so no non-existing objects need enter the picture. In order to meet the constraint that we give a unified account of the appearing relation, one would have to say that even in the case of veridical perception the relation of appearing holds between a subject and a mental representation.⁸

However, appeal to a mental representation of an object as opposed to the object itself might seem to merely push the problem back a step, onto the representation itself. We start out wondering what it is to have a merely intentional object appear to a subject, so we substitute a representation of a non-existent object to take care of the apparent commitment to non-existent entities. But now the question immediately arises, what is it to represent a non-existent object? Again, with Santa Claus, it's clear how the representation of Santa Claus exists, in contrast with Santa Claus himself, but then we want to know what makes this particular representation "of Santa Claus"? It looks as though we need to appeal to a merely intentional object again, and we've made no progress.

There's another problem with the mental-representation move as well. If one literally substitutes representations of objects for the objects themselves in the appearing relation, we seem to be committed to the claim that what appears to us are our own ideas, not the objects the ideas are about. While ultimately we may be committed to something in the vicinity of this claim—and it may not be so hard to live with either—clearly the claim can't be right, baldly stated in this way. After all, I do see red tomatoes, not ideas of red tomatoes. So some more complicated story about how the appearing relation should be analyzed, along with a subtler account of the relation between the appearing relation and standard cases of seeing objects is needed.

To understate wildly, I don't have a fully worked out account to offer. But here is the picture I'd like to try on.⁹ I'm tempted to think that Kant, as I will update him and interpret him (and I make no claim to historical

accuracy here), basically got it right about the world as we perceive it. According to Kant, objects we perceive are "appearances," not things in themselves. For him this followed from the fact that Space and Time are forms of intuition, imposed on the framework of the external world by the mind. I am not making that claim in particular. But what I do want to say is Kantian in spirit. I want to say that, through its causal interaction with the world, the mind constructs a virtual world. This world is then the immediate arena of experience, and its objects are the intentional objects we're looking for as the other relata of the appearance-relation. When this construction results from the right kind of interaction with the world as it really is, then the relevant perceptions are veridical. When this construction is caused internally without the right kind of connection to what's outside, it's hallucinatory. But the point is that when we are appeared to by a red tomato, the immediate relatum of the appearance relation, what it is that appears, is a virtual—or intentional—object; or, as Kant would have it, an "appearance."

The problem with intentional objects, as I see it, is that they have a definite role to play in psychological explanation, but it's one that is hard to play without being real. Most of our psychological states seem to be directed on objects, and to properly characterize them it seems necessary to advert to what they're about—their contents—and not merely their internal, intrinsic properties. Leaving out the contents, including the objects the states are about, seems to leave out what makes them psychological, and not merely formal. But unless the objects are real, it's difficult to see how they can play the requisite role.

So suppose we take seriously the idea of a virtual world as an explanatory posit. It turns out, on this view, that the only way to properly understand the mind is to characterize it as a kind of story-teller, and to take the world it creates for itself seriously in the characterization of its states. Again, echoing Kant, the world as experienced is a world of "appearances." But of course this isn't pure fiction. The external world, the world described by physics and the other natural sciences, constrains the construction of the world of appearances in such a way that it makes sense to speak of some appearances as accurate and others as illusory, and still others as flat-out hallucinatory.

One reason for adopting such a view anyway is the mess we get into when we try to construct a uniform metaphysical picture that answers both to the demands of natural science and to those of psychology. Questions

of endurance vs. perdurance, the reality of macro objects, and other metaphysical puzzles of this sort arise because of the clash between, as Sellars (1963) put it, the manifest and scientific images. Physics tells us there are basic entities like quarks and electrons doing their thing, and has no need of tables and chairs. But you really can't explain our behavior without taking seriously that we live in a world of tables and chairs. So the Kantian view, that the physical world of quarks causally interacts with our minds in such a way as to produce a reliable construction of a world of tables and chairs makes sense. The tables and chairs are there in the world as we experience and think about it, but not in any metaphysically serious way that would require some account of precisely which collections of quarks a chair is.

Leaving aside the question of any independent motivation for positing a Kantian virtual world, how does this idea answer to the initial motivation behind the AA model? Aren't we now back to locating phenomenal character and secondary qualities inside the mind? Of course in a way that's right. But remember the problem was to respect transparency while doing that. We didn't want phenomenal character to turn out to be experienced as just a property of oneself. It needed to be characterized as essentially a mode of presenting the world.

As I understand the AA model now, this spirit of transparency is preserved. The mind, as a causal result of its interaction with the physical world,¹⁰ goes into a state in which it constructs a representation of a world, to the intentional objects of which it stands in the AA relation. So when light bounces off that particular surface on the tomato on the counter and enters my eye, I go into a state of being-appeared-to-redly—a state that is properly characterized only relationally, treating the tomato-representation that is part of it as an object. Consciousness thus involves constructing its own objects, and phenomenal character is a property of how it is intentionally related to them. To say I really understand what I'm saying here would be overly optimistic. But I do think consciousness is a mystery, that intentionality is part of that mystery, and our qualitative experience, as the locus of that mystery, demands we look for a new (or maybe an old) model to understand it.

Joseph Levine

University of Massachusetts
Amherst

NOTES

1. What I'm calling the 'AA model' here very closely resembles what has traditionally been called the 'theory of appearing', which I only discovered after arriving at the view through my own route, which I trace in this paper. See Langsam (1997) for a defense of the view. I differ from Langsam in a crucial respect, because he endorses disjunctivism, which I reject in Section 4.
2. Shoemaker (2001/2002), whose view I will focus on below, emphasizes the importance of color constancy to this topic.
3. Of course it's not that "full possessor" of a concept is itself a well-understood notion.
4. Campbell (forthcoming) emphasizes this point, though he makes different use of it.
5. Though it's an interesting sociological note that whenever this question is raised at conferences including both philosophers and psychologists, the latter almost invariably support irrationalism. So I suppose the phrase 'prima facie implausibility' must be relativized to discipline.
6. The idea that perceptual experience presents two properties, objective ones and appearance properties, is developed in Shoemaker (2001/2002). In Shoemaker (manuscript) he takes back the idea that an appearance property is represented, but something quite analogous is still posited. There isn't space here to explore the subtlety and complexity of Shoemaker's different positions on this issue.
7. Indeed Byrne (2001) includes sense-data theory as a version of the intentionalism about qualia he's defending.
8. Does this mean the AA model collapses back into the representation model? No, for the appearing relation is still kept intact as a primitive, conscious, intentional relation. It's just that one of its relata is now a mental entity. "Curiouser and curiouser," as David Lewis might say.
9. I am indebted in my thoughts about this, though in ways he wouldn't anticipate, to many conversations with Georges Rey, who has been attempting to develop a view of intentionality involving a metaphysically harmless quantification over nonexistent objects (see Rey forthcoming). Of course he wants to rid the world of qualia this way, whereas I embed them as properties of the relation we bear to the intentional objects. Exploration of the difference between our views and why I stubbornly hold on to the reality of qualia in this way will have to await another paper.
10. How? Well that's the physical-realization problem, and I don't have an answer to that, and suspect there isn't one.

REFERENCES

- Balog, K. (1999). "Conceivability, Possibility, and the Mind-Body Problem," *The Philosophical Review*, 108: 497–528.
- Byrne, A. (2001). "Intentionalism Defended," *The Philosophical Review*, 110: 199–240.
- Campbell, J. (forthcoming). "Transparency vs. Revelation in Color Perception," *Philosophical Topics*.
- Dretske, F. (1995). *Naturalizing the Mind*. Cambridge, MA: Bradford Books/M.I.T. Press.
- Hardin, C. L. (1988). *Color for Philosophers: Unweaving the Rainbow*. Indianapolis, IN: Hackett.

- Harman, G. (1990). "The Intrinsic Quality of Experience," in J. Tomberlin, ed., *Philosophical Perspectives, 4, Action Theory and Philosophy of Mind*. Atascadero, CA: Ridgeview, 31–52.
- Hilbert, D. R. (1987). *Color and Color Perception: A Study in Anthropocentric Realism*. Stanford, CA: Stanford University Center for the Study of Language and Information.
- Langsam, H. (1997). "The Theory of Appearing Defended," *Philosophical Studies*.
- Levin, J. (2007). "What Is a Phenomenal Concept?," in *Phenomenal Concepts and Phenomenal Knowledge*, eds. Torin Alter and Sven Walter. Oxford: Oxford University Press.
- Levine, J. (1988). "Absent and Inverted Qualia Revisited," *Mind & Language*, 3, 4, 271–87.
- _____. (2003). "Experience and Representation," in *Consciousness: New Essays*, Quentin Smith and Aleksandar Jokic, eds. Oxford: Oxford University Press.
- _____. (2007). "Phenomenal Concepts and the Materialist Constraint," in *Phenomenal Concepts and Phenomenal Knowledge*, eds. Torin Alter and Sven Walter. Oxford: Oxford University Press.
- Loar, B. (1997). "Phenomenal States," in Block, N., Flanagan, O., and Güzeldere, G. eds. *The Nature of Consciousness*. Cambridge, MA: M.I.T. Press.
- Lycan, W. G. (1996). *Consciousness and Experience*. Cambridge, MA: Bradford Books/M.I.T. Press.
- Martin, M. G. F. (2004). "The Limits of Self-Awareness," *Philosophical Studies*, 120: 37–89.
- Papineau, D. (2002). *Thinking About Consciousness*. Oxford: Oxford University Press.
- Perry, J. (2001). *Knowledge, Possibility, and Consciousness*. Cambridge, MA: M.I.T. Press.
- Rey, G. (1998). "A Narrow Representationalist Account of Qualitative Experience," in Tomberlin, J. ed., *Philosophical Perspectives 12: Language, Mind, and Ontology*, 435–57, Oxford: Blackwell.
- Rey, G. (forthcoming). "Phenomenal Content and the Richness and Determinacy of Experience," *Journal of Consciousness Studies*.
- Rosenthal, D. (1997). "A Theory of Consciousness," in Block, N., Flanagan, O., and Güzeldere, G. eds., *The Nature of Consciousness: Philosophical Debates*, Cambridge, MA: M.I.T. Press, 729–53.
- Sellars, W. (1963). *Science, Perception, and Reality*. London: Routledge and Kegan Paul.
- Shoemaker, S. (2001/2002). "Introspection and Phenomenal Character," *Philosophical Topics*, reprinted in Chalmers, D. ed., *Philosophy of Mind: Classical and Contemporary Readings*. New York: Oxford University Press.
- Shoemaker, S. (manuscript). "On the Ways Things Appear."
- Tye, M. (2000). *Consciousness, Color, and Content*. Cambridge, MA: Bradford Books/M.I.T. Press.