

Color, Externalism, and Switch Cases

David Bain
University of Glasgow

Abstract

I defend externalism about color experiences and color thoughts, which I argue color objectivism requires. Externalists face the following question: would a subject's wearing inverting lenses eventually change the color content of, for instance, those visual experiences the subject reports with "red"? From the work of Ned Block, David Velleman, Paul Boghossian, Michael Tye, and Fiona Macpherson, I extract problems facing those who answer "Yes" and problems facing those who answer "No." I show how these problems can be overcome, leaving externalism available to the color objectivist.

I defend in this paper a view about color by defending something I argue the view requires: an externalist conception of both color experiences and color thoughts (e.g., beliefs and judgments). The view is that colors are objective, real, and visually represented, sometimes veridically. Call this the ORV account. One version says that colors are reducible to microphysical properties (whether categorical or dispositional); another says they are primitive, at most supervening on the microphysical. Either way, they are objective, or mind-independent, in the sense that the mental is not mentioned by a complete account of the colors. ORV further takes them to be real, in the sense of actually being possessed by material objects. By contrast, some error theorists think colors are phenomenal properties of experience, "projected" onto the perceived world and hence

David Bain is a lecturer in philosophy at the University of Glasgow. His recent research has focussed on philosophy of mind, particularly perception and bodily awareness. Publications include "Intentionalism and Pain" (Philosophical Quarterly, 2003), "Private Languages and Private Theorists" (Philosophical Quarterly, 2004), and "The Location of Pains" (Philosophical Papers, 2007).

are neither objective nor real; and some dispositionalists think they are dispositions to cause experiences, hence real but subjective.

ORV faces a difficulty whose solution, I argue in section 1, requires an externalist conception of both color experiences and color thoughts. Section 2 poses a question for this conception: if a subject underwent a context switch, involving her wearing inverting lenses over an extended period, would she undergo a change with respect to the color contents of visual experiences of given types? Call this the visual change question, and call accounts Yes views or No views depending on their answer. Since Ned Block raised this question in 1990, a number of objections have been made against Yes views and, very recently, a challenge has been posed for externalist No views. Here I reply to the former (sections 3–8), and also show how to meet the latter (sections 8–9). Both Yes views and No views, I conclude, remain viable options for the externalist. This means we can let the visual change question be settled by, rather than determine, our choice of externalist psychosemantics. It is also good news for ORV.

Some terminological preliminaries: I speak of an experience's *representational content* (or just *content*) and *phenomenal character* (or just *character*). "Content" refers to what the experience represents as being the case (e.g., that there is a red tomato in front of you); "character" refers to what it is like for the subject to undergo the experience. Here "red-representing" refers to the possession of a particular color content, that is, the representation of a color, in this case redness.¹ I take "The tomato looks red to me" to mean that my experience of the tomato *represents* it as being red. Not only experiences but thoughts have color content, but I use "*visual* color content" to refer specifically to the content of experiences. "Red-feeling" refers to a particular color character, that is, the character (if any) of the visual experiences humans normally undergo when viewing red things. "Qualia" refers to *nonrepresentational* phenomenal properties of experience (if such exist); "red*" refers specifically to the quale (if any) shared by the experiences in which we normally see red things.

With these preliminaries completed, let us turn to a difficulty ORV faces.

1. The Symmetry Problem

To appreciate a worry for ORV, assume ORV is true and that in normal humans and normal circumstances red things cause *r*-representing experiences and green things *g*-representing, where *r* and *g* are colors at the same level on the determinable-determinate hierarchy—hence, neither is a shade of the other. (They might be perfectly ordinary colors, for example, red and

green; but I use “*r*” and “*g*” so as to remain neutral for now about which colors they are.) Imagine now another species, L-humans. L-humans are very like normal humans: they apply the same color terms (“red,” “green,” etc.) to the same objects as we, thereby effecting discriminations as fine-grained as ours; their nonverbal responses to colors are also the same (at least when nonintentionally characterized); and they evolved in the same lighting conditions as humans, surrounded by the same colors, albeit in a hitherto undiscovered region of the Earth. They are also exactly like us physically *except* they have one-millimeter thicker corneas, whose extra layers happen to be atom-for-atom identical to inverting lenses, that is, contact lenses that invert normal humans’ color experiences, at least when first worn—so that, for example, red things cause *g*-representing experiences and green things cause *r*-representing.² Suppose, finally, that Norm, a normal human, and Len, a normal L-human, are both viewing a red cube in broad daylight.

The worry is that Norm’s cube experience will be *r*-representing but Len’s will be *g*-representing. Yet the cube has only one color at the relevant level (redness), so *at most* one of Norm and Len can be seeing it veridically. Since it is hard to imagine what more veridicality could require, moreover, surely *at least* one of them is doing so. But which one? Call this the veridicality question. Given the perfect symmetry between Norm’s and Len’s relations to the world, there seem to be no possible grounds for selecting either of them in answer to the question. Thus it begins to look as though ORV is forced to treat as genuine a pseudo-question.³ This is the symmetry problem.

The problem could be overcome by rejecting ORV. One approach is to deny that exactly one of Norm and Len represents the cube veridically, while conceding they represent it as different colors. If (contra ORV’s realism) no objects are colored, for example, then neither Norm nor Len sees the cube veridically. Or perhaps both do, if (contra ORV’s objectivism) *r* and *g* are sensory dispositions that a single surface could coconstitute, for example, to cause red* experiences in humans, and to cause green* experiences in L-humans. Another approach is to deny that visual experiences have color content, in which case—just as when eating cabbage makes Norm but not Len nauseous—no question of correctness arises.⁴

But what if ORV is not rejected? ORV cannot deny that experiences have color content, by definition, but I think it should deny that Norm’s and Len’s cube experiences *differ* in color content, that is, deny that Len’s cube experiences could be *g*-representing if Norm’s are *r*-representing. If color character were independent of color content, ORV could still allow that Norm and Len see the cube differently, their experi-

ences sharing color content but differing in qualia; but I recommend below the strong representationalist view that color content constitutes color character. Hence I think ORV must deny that, with respect to color, Norm and Len see the cube differently at all.

One objection to this is that different species obviously see the world differently. Another objection considers a normal human who has been wearing inverting lenses only momentarily—call him Abe. Abe sees the cube differently from Norm, thanks to the lenses. But suppose we stipulate that Abe and Len are *virtually* internally identical when viewing the cube, that is, atom-for-atom the same except that Len has an extra corneal layer where Abe has lenses.⁵ From this, one might infer that Abe and Len see the cube identically to one another. It follows by transitivity that Norm and Len see it differently from one another.

But, in fact, neither objection establishes that Norm and Len see the cube differently. With respect to color, humans do indeed visually represent objects differently from other species, but the differences arguably consist in asymmetries that ground acceptable answers to the veridicality question, asymmetries that do not exist between humans and L-humans. As for the transitivity objection involving Abe, ORV should embrace the following crucial claim:

E-externalism

The color contents of visual experience are broad, that is, they do not supervene on the subject's internal properties.

Conjoined with strong representationalism, defended below, this entails that color character too is broad. Hence one might reply to the transitivity objection that even if Abe and Len were *perfectly* identical, internally, their cube experiences might yet differ in color content (hence, given representationalism, character), and so Len's and Norm's might yet share theirs. This reply is one route to ORV's commitment to e-externalism.

More generally, the motivation for e-externalism is that the only way ORV can prevent unanswerable veridicality questions arising in cases like Norm and Len's is by holding that the subjects' experiences in those cases share their color content, which in turn requires according matters beyond the subjects' skins a role in determining such content, hence e-externalism.⁶ This paper is largely a defense of e-externalism.

Now, more than one psychosemantics entails e-externalism. But here I want neither to endorse a particular psychosemantics for experience nor to rule out the idea that specifying necessary and sufficient conditions for this or that perceptual content is prevented by the uncodifiability of content-determining

relations. The point remains that precisely what made it impossible to single out either Len or Norm in reply to the veridicality question, namely the perfect symmetry between their relations to the world, makes plausible their experiences' representational sameness, thereby calling into doubt the question's presupposition: that just one of them is right.

Despite my neutrality, however, it will be useful in what follows to sketch three externalist psychosemantics. Call psychosemantics for visual color content *e-semantics* and psychosemantics for the color content of thoughts *t-semantics*. The three e-semantics I have in mind arise by replacing *C* in the following schema with the substitutions below:

S's experiences of type *X* represent a given property, *F*, iff: if *C* were to obtain, *X*-type experiences would occur in *S* iff, and because, an *F* was in view (or, for short: *X*-type experiences would *covary* with *F*).

Typical e-semantics: perceptual circumstances that are statistically typical for *S* when *S* undergoes the experience.

Teleological e-semantics: perceptual circumstances for which the visual system of *S*'s species evolved.

Optimal e-semantics: perceptual circumstances that are *optimal* in Michael Tye's sense, that is, (roughly) circumstances that have not been interfered with (Tye 2000, 136–39).⁷

Block (1990) and Tye (2000) respectively adopt something like the first and third e-semantics;⁸ Fred Dretske (1995), Bill Lycan (1996), and Ruth Millikan (1984) adopt something like the second. (The differences between their formulations and mine need not concern us here.)

Congenially, all three e-semantics entail that Norm's and Len's cube experiences share their color content. Suppose that in normal humans, in broad daylight, neural states of types *R* and *G* covary with red and green things respectively, the converse being true in L-humans (and humans wearing lenses). When viewing the red cube, therefore, Norm and Len are in different neural states: an *R* state and *G* state, respectively. Even so, each is in a state that *in him covaries with red things in the relevant circumstances*, in whichever of the three ways above those circumstances are characterized. So on each e-semantics, both Norm's and Len's cube experiences are red-representing. Not so Abe's, since he is in a state, *G*, that in him (and Norm) covaries—in the relevant circumstances—with green things.

Notice finally that we ought to embrace not only e-externalism but the following:

T-externalism

The color contents of *thoughts* are broad.

Further motivations for t-externalism emerge below, but one motivation echoes the foregoing: unless we embrace t-externalism, ORV will generate cases where two subjects *judge* an object with a given color to be different colors, and where it is an unanswerable pseudo-question which one is correct. In the final section, I recommend a particular externalist t-semantics.

The bottom line is that ORV theorists ought to be externalists. Let us turn now to a question that generates difficulties for e-externalism.

2. The Visual Change Question

Suppose inverting lenses are inserted into the eyes of a normal human subject, call her Traveller, at just the moment she is transferred from Earth to Inverted Earth, on which everything (including Traveller's body) has the color complementary to its color on Earth, so ripe tomatoes (hereafter, just "tomatoes") are green and grass is red (Block 1990). Changing the distal causes of her experiences, the insertion of the lenses is what is sometimes called a context switch. But Traveller remains unaware of both the switch and the transfer since their effects cancel out, at least at first: tomatoes, although green, still look red. Inverted Earthlings, moreover, speak a language identical to English but for an inverted color vocabulary, so they will say to Traveller, "Tomatoes are red," meaning that tomatoes are green. Suppose finally that Traveller sees a tomato in broad daylight seconds before moving to Inverted Earth (t_1), seconds after (t_2), and again 50 years later (t_3). This raises the visual change question: granted that at t_1 Traveller's tomato experiences—or, more generally, the experiences she reports using "red"—are red-representing, does their color content change by t_3 ?

In itself, e-externalism compels neither "Yes" nor "No" as an answer. Whereas Typical e-semantics generate Yes views, for example, the other two generate No views. (As I shall also put it, Typical e-semantics is a *switch* e-semantics; the other two are not.) Suppose, for example, the inverting lenses work as before; hence R states are caused in Traveller not only by the red tomato at t_1 but, given the lenses, by the green tomatoes at t_2 and t_3 . With which color would Traveller's R states covary in the circumstances relevant at each time? For Typical e-semantics, the relevant circumstances are those statistically typical for Traveller when undergoing the experiences in question. Hence they change: they include at t_3 but not at t_1 her wearing lenses, since at t_1 she has never worn lenses but by t_3 has done so for

50 years. Now Traveller's R states would covary with red things in circumstances typical at t_1 but with green things in circumstances typical at t_3 . Consequently, according to Typical e-semantics, her R states are red-representing at the earlier time and green-representing at the later. On Teleological and Optimal e-semantics, by contrast, they are red-representing throughout, since the relevant circumstances remain constant.

So, depending on their psychosemantics, e-externalists can return either answer to the visual change question. But whichever answer they give, difficulties arise, as we shall see presently. Yes views, ostensibly the more problematic option, occupy us for much of the paper. I consider in sections 3–4 Block's argument that they entail both color qualia and nonphenomenal visual color content. Against Block, I argue that it would be problematic if they did so, but also that they do not, and that Yes theorists can and should be representationalists. Representationalism, admittedly, can appear at least to aggravate three further difficulties I consider in sections 5–7: that Yes views are incompatible with introspective knowledge of current experiences (Boghossian and Velleman 1991); that there is no plausible account of *gradual* changes in visual content (Tye 2000; Macpherson 2005); and that changing memory contents cannot conceal *sudden* visual changes (Macpherson 2005). Notwithstanding representationalism, I argue all these problems can be overcome. I turn in section 8 to some final challenges, not only for Yes views but also for No views. These too, I claim, can be met, provided we embrace the account of color concepts I sketch in section 9.

3. Qualia and Nonphenomenal Color Content

Block is a Yes theorist. He thinks Traveller's lenses eventually induce a change in the color content of her tomato experiences. But what about those experiences' character? This, Block argues, remains constant. If he is right, we shall see below, Yes views entail color qualia and the nonphenomenality of visual color content. Whereas Block welcomes these putative consequences, I argue in the next section that they would make Yes views untenable.

Before turning to Block's argument, it will be useful to have before us a table of color judgments to refer back to:⁹

First-Order Judgments	Second-Order Judgments
<i>Present Tense Noncomparative</i>	
A ₁ . This tomato <i>is</i> red	A ₂ . This tomato <i>looks</i> red
<i>Past Tense Noncomparative</i>	
B ₁ . That tomato <i>was</i> red	B ₂ . That tomato <i>looked</i> red
<i>Diachronic Comparative</i>	
	C ₂ . This tomato looks the same way as that tomato looked
	D ₂ . The tomato experience I am now having has the same <i>color character</i> as the tomato experience I had earlier
E ₁ . This tomato <i>is</i> the same color as that tomato was	E ₂ . This tomato <i>looks</i> the same <i>color</i> as that tomato looked
	F ₂ . The tomato experience I am now having has the same <i>color content</i> as the tomato experience I had earlier

Now the reason Block thinks the color character of Traveller's tomato experiences remains constant is that she notices no change in them. In particular, on the basis of introspection and memory, she *claims* there has been no change. Consider her counterfactual counterpart, Untravelled, who stays on Earth but remains internally identical to Traveller. Suppose she says at t_3 , "This tomato looks the same way as that tomato looked" (C₂). Verbal dispositions being narrow, Traveller utters the same sentence and intends it, Block claims, as a judgment about character.¹⁰ If the character of Traveller's tomato experiences has changed, this judgment is mistaken; but Traveller has not made a mistake, Block thinks; hence, the character has not changed. Notice, by the same token, a second conclusion seems to follow: that Traveller's t_1 grass experience and t_3 tomato experience *differ* in color character. Traveller judges they differ, after all. And, given the phenomenal difference between her tomato and grass experiences at t_1 , the difference is also entailed by the conclusion that her tomato experiences' character is constant.

Putting together these conclusions about character, on the one hand, with a Yes view, on the other, the upshot is that Traveller's t_1 and t_3 tomato experiences differ in color content while sharing their color character (call this content indepen-

dence), and that her t_1 grass and t_3 tomato experiences share their color content while differing in color character (call this character independence).¹¹ Given Block's conclusions, therefore, Yes views "split" content and character, being inconsistent with the following:

Strong Color Representationalism

All color character is visual color content; hence, necessarily, two experiences differ in color character iff they differ in color content.

Weak Color Representationalism

All color character supervenes on (is determined by) visual color content; again, necessarily, if two experiences differ in color character, they differ in color content.

Both character and content independence contradict the strong thesis; character independence also contradicts the weak. Block welcomes these antirepresentationalist consequences and embraces qualia, that is, phenomenal properties that are neither determined by nor identical to representational properties.

He also goes further. Notice, after all, one might reject color representationalism while accepting the following:

Mixed View

Some color character is (alternatively: supervenes on) color content; again, there are aspects of color character such that, necessarily, two experiences differ in respect of those aspects iff (alternatively: only if) they differ in color content.

An elaboration of a case considered by Christopher Peacocke illustrates the view (1983, 39). Consider the experiences undergone when seeing (i) an array of objects including a piece of white paper, in normal conditions, (ii) the same array but with green rather than white paper, (iii) the first array, but viewed through red glass. The three experiences all differ in color character. But a mixed theorist might say that the first two differ in content rather than nonrepresentational features, whereas the first and third must differ nonrepresentationally, since the paper is represented the same way (namely, as white) in both. The view claims that, given the latter, there are color qualia, but given the former, color content is nonetheless phenomenal.¹²

Whatever this view's merits, it is not Block's view. He takes color content to be nonphenomenal: no differences in color character are constituted by, or entailed by, differences in color

content. For any color experience, after all, Block appears to think there is a possible experience sharing its character but lacking content altogether (1996, 26, 28). Hence, for any aspect of color character, two experiences might differ in respect of it while sharing their representational blankness. Moreover, if any differences in color character were to be determined by color content, plausibly they would include the difference in color character between one's experience in broad daylight of red tomatoes, on the one hand, and of green grass, on the other. But if Block is right, *that* phenomenal difference is *not* determined by color content, for that experience of green grass, on the one hand, and an experience of an Inverted Earth tomato 50 years after having inverting lenses inserted, on the other, would exhibit the very same phenomenal difference while sharing color content.

So if Block's argument is sound, Yes theorists are committed to qualia and the nonphenomenality of visual color content. Block embraces these commitments. I shall now argue he should not.

4. The Elusiveness Objections and the Memory Strategy

A complete case against color qualia and nonphenomenal color content is beyond this paper's scope, but I want at least to gesture at why I think these putative commitments seriously disfigure the notions both of character and of content. However, I shall also argue that Block is wrong to think they are commitments that Yes views need to incur.

Start with character. When first introduced in terms of the "what it is like" locution, character is apt to strike us as a familiar, introspectible, reportable feature of consciousness. But this conception comes under pressure when character and content are split.

For one thing, Block thinks qualia might vary extensively amongst normal humans. So if "red," "green," and the rest of our ordinary color vocabulary could be used (e.g., in "looks"-constructions) to report qualia, then that vocabulary's meaning would have to vary extensively amongst us, which Block finds highly implausible. Hence, Block thinks (rightly, I suggest) that when we say something is or looks *red* (A_1, A_2), or that two things are or look the same *color* (E_1, E_2), we are endorsing or reporting visual color *content*. We are not reporting qualia. But then, given Block's view that character is constituted by qualia alone, we are also not—rather surprisingly—reporting character. Block is unequivocal: "[p]henomenal character," he says, "is not expressible in English" (1996, 47; see also 1990, 55–57). Admittedly, he allows that one use of "This thing looks the same way as that thing" (C_2) expresses a judgment about character,

but there still is no shared vocabulary in which to express *which* way the two things look (in that sense), which character the two experiences share. “Red*” is not a counterexample since its reference was supposed to be fixed to the quale of experiences humans normally undergo when seeing red things, whereas this description may well fail to denote if there is extensive variation in the qualia humans enjoy.¹³ Hence, at best, the reference of “red*” and its ilk must instead be fixed in our own cases. So if qualia exist, they are, if you like, semantically elusive. And if character is constituted by qualia alone, as Block thinks, character too is elusive.¹⁴ Now, representationalism, by contrast, generates no such surprises. It straightforwardly allows what is plausible: that color character is what you express when, for example, you say that a thing looks red to you.

Qualia’s elusiveness, moreover, is arguably epistemic as much as semantic. At any rate, the most compelling account of introspective knowledge of current experiences does nothing to explain knowledge of qualia. The account is Gareth Evans’s (1982, 227–30). To find out what color something looks to you, he claims, you determine which perceptual judgment you would—if deprived of collateral information—make about its color on the basis of your visual experience. Then you preface the result with “It visually seems to me as though.” For example, “It visually seems to me as though the tomato is red” or, equivalently, “The tomato looks red to me.” It seems very plausible that this procedure is at the heart of our knowledge of our current experience, but notice it will yield knowledge only of content, not of qualia, since it is content that our perceptual judgments endorse. So if qualia alone constitute character, we face the unenviable task of otherwise explaining how Traveller comes to know, using introspection and memory, that her earlier and current tomato experiences share their character, and even how Traveller comes to know what (incommunicable!) character her current tomato experiences have.¹⁵

Ironically, then, one cost of saving Traveller from error about character by splitting character and content is that character becomes a highly elusive feature of consciousness. So, I now claim, does visual color content.

To see this, it is first useful to raise a question about Block’s antirepresentationalist argument. Notice that, using Untravelled, we can straightforwardly motivate the idea that Traveller claims that her t_1 and t_3 tomato experiences share their color *content* (she might utter, after all, E_2 , F_2 , and, intending it as a judgment about content, C_2) even though on Block’s view they do not. Hence, Block must allow that there is reason to convict Traveller of a diachronic error (a mistaken diachronic comparison) about *content*. Yet his antirepresentationalist argument involves *denying* that Traveller thereby makes a diachronic mistake about her experiences’ *character*. So a crucial question

arises: what non-question-begging reason can Block have for this denial? Again, what justifies the difference in his attitude toward mistakes about content and mistakes about character?

One answer is suggested by Block's rhetorical question, "Why shouldn't we believe [Traveller]?" (1996, 42). There simply is no *reason* to suspect a mistake about character, his idea seems to be. But that is question-begging. If color content is color character, as the representationalist claims, then Block's reasons for convicting Traveller of mistakes about the former are—if sound—reasons to convict her of mistakes about the latter. Considerations of charity also do not answer the question, for Traveller's erroneous judgment about content is not eliminated by claiming she additionally makes a correct judgment about qualia. Nor is the question answered by saying that diachronic mistakes about character are impossible, for surely they actually occur—when, for example, the difference in character is small, or the change slow, or the earlier experience misremembered.

Another answer, this time concerning explanation, is suggested by Block's claim that "it is a necessary feature of phenomenal character that if a change is big enough and happens fast enough, we can notice it" (1998, 668, emphasis removed) and that character "is what is relevant to noticing a difference" whereas content "can change purely externally" (1996, 47). This suggests the following line of thought: Mistakes about character need to be explained consistently with conceiving character as 'relevant to noticing a difference'. But, given that the phenomenal difference the representationalist alleges between Traveller's t_1 and t_3 tomato experiences is great, and that Traveller may well remember her earlier experience perfectly, there could be no such explanation of her making a mistake about character. By contrast, the idea goes, since relevance to noticing differences is *not* part of the idea of content, we need no similar explanation to make sense of a mistake about content. We need simply deny that it is thereby a mistake about character.

If this is Block's line of thought, we are now in a position to see what is wrong with his conception of visual color content. It is a conception on which such content is not itself "relevant to noticing a difference." Again, the idea is that such content is not itself a feature in virtue of which subjects can discriminate their color experiences or the colors of things. This is a highly problematic conception, rendering unavailable a crucial idea I recommend in the final section: that such content acquaints us with the colors, that it determines the color content of our thoughts, in particular that it enables us to meet the requirement that we *know which* objective properties are the referents of our color concepts, and that it enables us to meet this requirement *by* grounding capacities to discriminate the different colors (and thereby, if Evans is right, the experiences repre-

senting them). Visual color content cannot determine our color concepts, I suggest, if at best it is contingently associated with the features in virtue of which we discriminate colors and our experiences. In short, if Block's argument is right, there is a sense in which visual color content too is epistemically elusive—and problematically so.

I suggest instead that such content is itself relevant to noticing a difference, indeed that it is phenomenal. But that means we need an explanation rendering intelligible Traveller's *failure* to notice the difference between her earlier and later contents. Happily, there is such an explanation, one invoking memory. Given that memories are contentful states, switch externalists have reason to think that by t_3 Traveller will *misremember* her t_1 experience's content.¹⁶ Block himself is a switch externalist. Indeed, he holds not only a switch e-semantics, as we have seen, but a switch t-semantics, on which Traveller's thought contents are sensitive to the practice of her linguistic community. Hence, he thinks that, when her linguistic allegiances shift from the Earthling community to the Twin Earthling community, she comes to use "red" to mean green. Her thoughts consequently change: those she expresses with "red" come to concern greenness. So, whether Block thinks of Traveller's memories of her earlier experiences as themselves experiences or thoughts, even he has reason to think that by t_3 they have undergone a content change, rendering them inaccurate. Consider the following: when at t_3 Untravelled remembers the t_1 tomato looking red, how does Traveller remember it looking? Well, she expresses the memory using "red," meaning green; and the experiences of hers that were once red-representing are now green-representing. Hence, surely she remembers it—or, rather, *misremembers* it—looking green. Given that by t_3 Traveller misremembers her t_1 experience's content, we can (as we hoped) reconcile her diachronic error about content with such content's being relevant to noticing a difference, indeed being phenomenal.

Notice, crucially, that this strategy undermines Block's anti-representationalist argument. What has in effect been shown, as Lycan too points out (1996, 129–34), is that we possess what Block's antirepresentationalist argument assumes we lack: an explanation making sense of a diachronic error on Traveller's part about the character of her own experiences. Call this reply to Block's argument the *memory strategy*.

Block (1996, 44) appears to suggest that the strategy is to infer a failure to *remember* character from the assumption that Traveller has made a *mistake* about character. But in fact the strategy is to infer the memory failure from representationalism and switch externalism about memory contents. Does this assumption of representationalism render the strategy question-begging, as Block (1996, 44–45) also appears to think? Surely

not. Block's argument in effect challenges the representationalist thus: if representationalism were true, what explanation would make sense of Traveller's consequent error about character? Thus challenged, it is perfectly legitimate to reply by giving the explanation that would be available *if* representationalism were true. Block (1996, 45) adds that even if Traveller were herself a switch externalist, she would have no reason to withdraw her claim that her t_1 and t_3 tomato experiences share their color character. But the point surely is that she would *if* she were also a representationalist.

Let us take stock. Yes theorists, I have suggested, need to do two things: deny nonphenomenal visual color content and embrace representationalism, thus rejecting qualia. Thanks to the memory strategy, they can do both.

5. The Introspective Knowledge Objections

At roughly the same time as Block was originally deriving antirepresentationalist conclusions by applying his switch e-semantics to Traveller's case, Paul Boghossian and David Velleman were also applying such e-semantics to switch cases, but doing so instead to provide a *reductio* of the semantics (1991, 89–90).¹⁷ They point out that if Yes views were correct, then a victim of a context switch would eventually end up committing a diachronic mistake about how things look, for example, falsely judging that tomatoes look the same color as they used to.¹⁸ We have conceded as much; but Boghossian and Velleman think something absurd follows: that a switch victim could not introspectively know even what color a tomato *currently* looked to be (1991, 90). Indeed, they appear to suggest further that, being unable to rule out a switch, *nonvictims* too would lack such introspective knowledge. Boghossian and Velleman do little to explain either claim. Here I consider three explanations of the inference from Traveller's diachronic mistake to her ignorance of current content. None, I argue, sustains the inference. Absent some further argument, I conclude, Yes views entail introspective ignorance neither in Traveller's case nor, a fortiori, in ours.

I take it, incidentally, that the putative entailment of introspective ignorance ought to alarm even nonrepresentationalists. Even if content did not determine character, it would be intuitive that a switch victim knows what colors things currently look to her. Representationalism might nonetheless seem to aggravate the difficulties: if ignorance of current content seems bad, ignorance of current character might seem worse. But I shall suggest that representationalism is, if anything, a help rather than a hindrance when dealing with Boghossian and Velleman's worries.

Now, the first worry Boghossian and Velleman might have in mind is that at t_3 Traveller not only makes diachronic, comparative mistakes about her experiences (E_2); her noncomparative judgments about her *current* experiences (A_2) are also false. Call this the truth problem. To reply to the truth problem, let us first enrich our taxonomy of views as follows:

(i) Does the color content of Traveller's tomato experiences change between t_1 and t_3 ?	(ii) Does the color content of the thoughts Traveller expresses with "red" change?
Yes	Yes
Yes	No
No	Yes
No	No

Call a view a Yes/No view, for example, if it answers (i) "Yes" and (ii) "No." Now, the truth problem surely undermines Yes/No views and No/Yes views. On a Yes/No view, for example, although tomatoes look green to Traveller by t_3 , she still judges—falsely—that they look red. But Yes/Yes views fare better. Block's view illustrates the point. Recall he accepts both a switch e-externalism and a switch t-externalism; hence, he can claim that although the color content of Traveller's tomato experiences changes by t_3 , her judgments about her experiences also change, at least partly as a result of her shifting linguistic allegiances, and they change in such a way as to remain true. When, for example, she utters at t_3 , "The tomato looks red" (A_2), she thereby judges that it looks *green*, which Yes views say it does at t_3 . I discuss a complication for this Yes/Yes reply in section 8, but on its face it meets the threat to the truth of Traveller's judgments about current content.

A second worry concerns discrimination. Consider first another switch case, in which Sally, a subject with no knowledge of chemistry, is switched from Earth to Putnam's Twin Earth, where instead of water there is twater, a distinct but superficially identical liquid (Putnam 1975). Consider the thoughts Sally expresses by uttering "Water is wet." Before the switch, these concern water, but 50 years after the switch they concern twater, according to switch t-externalism (extended to the concept *water*). Hence, such t-externalism appears to threaten Sally's knowledge of her thoughts, much as switch e-externalism appeared to threaten Traveller's knowledge of her experiences. Expressed as a truth problem, however, the worry about Sally can be met by a strategy structurally similar to the Yes/Yes approach above: although Sally eventually thinks first-

order thoughts about twater rather than water, her introspective thoughts *about* her current first-order thoughts also change, and in such a way as to remain true. But Jessica Brown and others have raised a further worry (Brown 2004). Sally still falsely judges that she is thinking the same thoughts she used to, and this raises the concern that she lacks a discriminatory capacity required for knowledge. Alvin Goldman (1976), for example, famously emphasizes the connection between knowledge and discriminatory capacities: a true visual judgment that something is a barn counts as knowledge, he thinks, only if the subject can tell barns apart from relevant nonbarns.¹⁹ In that light, Sally's apparent inability to tell her twater thoughts apart from her earlier water thoughts can appear to threaten even her noncomparative knowledge of her current thoughts—for example, her knowledge that she is thinking that twater is wet.²⁰ A memory strategy might seem to overcome the worry: Sally's diachronic mistake, the idea goes, reveals not the absence of a discriminatory capacity required for knowledge, but her misremembering her earlier water thought as a twater thought. But Brown (2004, 56–58) counters that even if Sally does misremember, there is still no reason to think she has the ability to tell the two types of thought apart, and some reason to think she lacks it, for example, her describing the liquids that both types concern in just the same terms (e.g., “clear, tasteless, potable liquid”).

Now this dialectic regarding Sally suggests a new worry back in the case of Traveller's color experience: that Traveller lacks a discriminatory capacity required—in addition to truth and perhaps reliability—for knowledge. Call this the *discrimination problem*.

However serious it is in Sally's case,²¹ I think the discrimination problem can be overcome in Traveller's, provided we take visual color content to be phenomenal. There is a disanalogy between the two cases such that Brown's reply to the memory point in Sally's case lacks a parallel in Traveller's. There is excellent reason to think that Traveller's failure to tell apart her t_1 and t_3 tomato experiences reveals memory failure, or at any rate *not* an inability to tell apart red-representing and green-representing experiences. Whereas Sally arguably comes to have twater thoughts *instead* of water thoughts, and arguably could not tell them apart even if she had them simultaneously, Traveller continues to have both red-representing and green-representing experiences and is manifestly able to tell them apart. She can tell at t_3 , for example, that tomatoes and grass currently look different colors. Hence, if the difficulty Boghossian and Velleman are raising is the discrimination problem, it fails to touch those Yes/Yes views that take visual color content to be phenomenal.

Notice, as an aside, that Block might say that the phenomenality of color content is inessential to my reply, since on his

picture too Traveller can tell apart her t_3 red-representing and green-representing experiences. She does so in virtue of their different qualia. But Block does look more vulnerable to the discrimination problem than I. Recall Traveller *cannot* tell apart the contents of her t_1 and t_3 tomato experiences; and unlike me, Block must deny the reason is memory failure, even if that occurs. Rather, on his picture, the reason is that none of us can ever tell apart visual color contents per se, but only the qualia to which they usually correspond, and in Traveller's case the context switch has upset the usual isomorphism: despite sharing their qualia, her t_1 and t_3 tomato experiences differ in content. Indeed, on Block's picture, it seems, even if Traveller had such experiences simultaneously, she could not tell them apart. In reply, Block might deny that Traveller could have such experiences simultaneously. But rather than pursue this aside, notice we are now in the vicinity of a third worry that Boghossian and Velleman might have in mind.

The third worry is semantical, not purely epistemological, and takes us back to section 4: if Block's picture were right, and visual color content were itself nonphenomenal—or at least not relevant to noticing a difference—then such content could not play its crucial role in *acquainting* us with the colors, in grounding our color concepts. So construed, Boghossian and Velleman's *reductio* would amount to this: "If Yes views were correct, then at t_3 Traveller could not know that the tomato looks green, simply because she could not know *which* property greenness is; she could not possess the concept *green*. For Yes views entail that visual experiences' contents have changed between t_1 and t_3 (her red-feeling experiences representing redness then greenness, her green-feeling experiences representing greenness then redness), *despite* there being no phenomenal change, indeed no change in an experiential feature relevant to noticing differences. If that has happened in Traveller's case, she cannot know which property greenness is. Indeed, if that *could* happen, none of us can."

My reply is straightforward. Taken as an objection against Block, this is just the point I made in section 4. But taken as a *reductio* of Yes views *tout court*, it fails to touch those on which visual color content is phenomenal.²²

6. The Intermediate Contents Objection

Let us now turn to a dilemma for Yes views. Call Yes views on which the change in Traveller's visual color contents occurs slowly *gradualist*, and the rest *nongradualist*. If gradualism is true, Tye (2000, 134–35) and Fiona Macpherson (2005, 133–34, 149) argue, Yes views are undermined by what I call the *intermediate contents objection*; but if nongradualism is true,

Macpherson thinks (2005, 133–34), the memory strategy fails. In this section, I consider this dilemma’s first horn.

To reconstruct the objection, as I understand it, consider these schemas:

- S At t , the tomato looks C to Traveller
- S_r At t , the tomato looks *red* to Traveller
- S_g At t , the tomato looks *green* to Traveller

Suppose the content change in Traveller’s tomato experiences occurs slowly over some period between t_2 and t_3 , between which times she (wearing lenses, recall) views a green tomato every ten seconds. S_r is determinately true where t is t_2 , and S_g is determinately true where t is t_3 . But, given that the change is gradual, there are intervening times at which Traveller is viewing a tomato but neither S_r nor S_g is determinately true. Call all such intervening times t_i . Now, Tye and Macpherson’s objection can, I think, be presented as the following challenge for the gradualist: granted that, at a given t_i , it is determinately true neither that the tomato looks red nor that it looks green, what color is it determinately true that the tomato looks? Again, the idea appears to be that, since the substitution of both “red” and “green” for C fails to produce a determinately true instance of S for any t_i , the gradualist must provide for each such time some other color predicate whose substitution *does* produce a determinately true instance.

Some gradualists try to meet the challenge head on. Bill Brewer proposes that, between t_2 and t_3 , colors “seep out” of Traveller’s visual contents, and when she can least well discriminate red and green she rather represents things as “black, white, and shades of grey” (2001, 874). Gradualists might alternatively suggest the substitution of disjunctive predicates such as “red-or-green.” Tye rejects this disjunctive proposal since red and green (or blue and yellow, in his example) are opposites on the hue circle (2000, 135). But it is unclear why that is problematic. Macpherson rejects the disjunctive proposal because “there is no plausible candidate for the phenomenal character of experiences with such unusual content” (2005, 134). But for strong representationalists the contents *are* the character—unusual contents, certainly, constituting unusual character. Still, I too reject both gradualist proposals, since they fail to capture the smoothness of the change even if they are *prima facie* plausible as accounts of its mid-point.

But there is, I suggest, a different and neglected approach gradualists can take: namely, to question the legitimacy of the challenge. Certainly the challenge is legitimate on the assumption that, for each t_i , since neither S_r nor S_g is determinately true, both must be determinately false and some other instance

of S be determinately true instead. But that assumption can be questioned by invoking vagueness. Again, it can be denied that gradualism is committed to an instance of S being determinately true for each t_i . Instead, gradualists can say that although at each t_i neither S_r^i nor S_g^i is determinately true, nor are both determinately false. Thus g gradualists need not cast about for predicates to substitute for C . They might rather treat the transition between t_2 and t_3 as revealing that “looks red to Traveller” is a vague predicate for which each t_i provides a genuinely borderline case. The point, notice, is not that “red” and “green” are vague, nor that each shade represented at a given t_i is a shade that is borderline between two colors, for example, red and orange. The point is that there could be borderline cases regarding the *representation* of a nonborderline shade of red—or of a nonborderline shade of green. (The view might be that “looks red” inherits its vagueness from a term such as “typical” in the e-semantics.) This, I suggest, is a promising line for gradualism to take.

Even if this proposal rests on a contested view of vagueness, it usefully indicates that the problem of intermediate contents is not specific to switch e-externalism, but exhibits a general philosophical difficulty. Switch t-externalism (extended to the concept *water*) faces what is essentially the same difficulty, of making sense of Sally’s gradual transition from water to twater thoughts. More generally, all except nativists face the kindred difficulty of making sense of the gradual acquisition of intentional states by beings who started with none. More generally still, there is the question of making sense of such gradual transitions as that from hirsute man to bald man. Hence, if the problem of intermediate contents remains a headache, it is not one that afflicts gradualists specifically.

7. An Objection to Sudden Memory Changes

What about the dilemma’s nongradualist horn? Gradualists and nongradualists alike need the memory strategy. But Macpherson argues that if nongradualism is true then the strategy fails, for it requires “changes in the character of one’s *memory* experiences that one does not, indeed cannot, notice” (2005, 133–34, my emphasis).

We can reconstruct Macpherson’s point, as I understand it, using the following diagram:

		<i>Times</i>		
		t_α	t_β	t_γ
<i>Level</i>	3			green- representing experiential <i>memory of</i> <i>memory at t_β</i>
	2		red- representing <i>experiential</i> <i>memory of</i> <i>experience at t_α</i>	green- representing experiential <i>memory of</i> <i>experience at t_α</i>
	1	red- representing <i>experience</i>	blinks	green- representing <i>experience</i>

Suppose nongradualism is true and, at a time between t_2 and t_3 , Traveller's tomato experiences suddenly change from being red-representing to being green-representing. Times t_α , t_β , and t_γ fall between t_2 and t_3 . At t_α , a given tomato looks red to Traveller (level 1, in the diagram); at t_γ , a moment later, Traveller blinks while visually remembering seeing the tomato a moment ago (level 2); another moment on, at t_γ , she views another tomato, but owing to a change in color content occurring at that very moment, the tomato looks green (level 1).

Now, Traveller will mistakenly judge at t_γ that the t_γ tomato looks the same color as the t_α tomato. The memory strategy, recall, explains this mistake in terms of Traveller's misremembering her t_α experience. But Macpherson, I take it, thinks this only defers the problem: the strategy invokes an unnoticed change in experiential memory (level 2) to explain why a change in visual experience (level 1) goes unnoticed; but what explains why the change in *memory* goes unnoticed? It is no less a phenomenal change, after all, given that the strategy invokes experiential, rather than merely propositional, memories.

In reply, notice first a scope ambiguity in the notion of failing to notice a change. To put it in terms of judgment, there is a difference between failing to judge there has been a change (call this the *weak sense*) and falsely judging there has been no change (the *strong sense*). If the change in memories at level 2 is unnoticed merely in the weak sense—if, for example, Traveller simply lacks states at level 3 representing her level 2 memories—then it is not obvious we need any explanation of the failure to notice. The hierarchy of higher-order representations must come to an end somewhere, after all. If, on the other hand, the change

in level 2 memories is unnoticed in the strong sense, this can be explained by redeploying the memory strategy: Traveller judges that her level 2 memories are unchanged since she has a level 3 memory of her earlier level 2 memory, and that level 3 memory is incorrect. After all, the reasons for believing there is change in Traveller's level 2 memories apply also to her level 3 memories. For example, "red" (meaning green) will figure in their verbal expression too: "At t_β I was remembering a tomato's looking red [green] at t_α ." Of course, Macpherson's demand for explanation will just iterate if Traveller has at the same level an immediately preceding memory with a different content. But here, and at every level up the hierarchy, the same strategy can be invoked, until we reach a failure to notice only in the weak sense. So Macpherson's objection fails, and both horns of the dilemma seem playable.

8. Further Challenges: Yes/Yes and No/No

I now take up two final challenges, one for No views, to which we at last return in this section, the other for Yes views. Here I explain the challenges; in the final section, I show how to meet them.

Recall that Yes views need to be elaborated as Yes/Yes views to overcome the truth problem, that is, to explain how Traveller's judgments about her current experiences remain true. But not just any Yes/Yes view will do. To ensure Traveller's thought contents change, Block, recall, invokes an idea associated with Tyler Burge (1979) and Putnam (1975): that thought contents are individuated partly by the linguistic practices of the subject's community. This is why Block populates Inverted Earth with the speakers of an inverted language, the idea being that Traveller's thought contents will eventually change as her linguistic allegiances transfer to her new companions. It is not clear whether Block would elaborate this account in such a way as to be vulnerable to the following point. But the point is important to register: that difficulties loom for any view that takes sociolinguistic factors to bear on the color content of thoughts but not on the color content of experiences. For one thing, even if in Traveller's case such a view entails that *both* the color content of her thoughts and the color content of her experiences change, it is unclear the view could guarantee the changes will occur *simultaneously*. Simultaneity is important because, without it, Traveller will make errors about her current experiences during the intervening period of mismatch. Such an approach also risks being committed to a Yes/No view of a new case, identical to the original Inverted Earth case except that Traveller finds Inverted Earth unpopulated (so that no pressure comes to bear on her linguistic allegiances). Indeed, it perhaps even risks commitment

to a No/Yes view of yet another case, in which Untravelled (on Earth, recall, without lenses) fails to notice that her linguistic community has suddenly been replaced with the speakers of an inverted color language (so that her linguistic allegiances are brought under pressure).²³ In short, the challenge for Yes/Yes views is how to achieve what we might call a *tight* Yes/Yes view, one guaranteeing the simultaneity of changes in visual and thought content, and one avoiding Yes/No and No/Yes commitments in these two new cases.

There is a related challenge for No views, to which I now finally return. These deny that Traveller's visual color contents change. I have adopted in this paper a position of neutrality between Yes views and No views. Even if Yes views are viable, as I have argued, one might have other reasons for preferring a *nonswitch* e-semantics (e.g., Teleological e-semantics) and hence a No view. No/No views in particular have their attractions, for example, ensuring that Traveller's judgments about current experiences remain true, without having to deal with—for instance—the above worries about simultaneity. But externalist No views do face a challenge.

The challenge, nicely posed by Macpherson (2005), is how to avoid a No/Yes view. Whether they are Yes theorists or No theorists, e-externalists tend to endorse t-externalism too. (The symmetry problem is one good reason why.) The difficulty arises when we examine *which* t-externalism No theorists endorse. Regarding thought content, even its color content, some No theorists are sympathetic to Burgean or Putnamian considerations about the relevance of the thinker's linguistic community and current environment. Hence, they endorse switch versions of t-externalism. Dretske (1995) and Tye (2000) are examples, combining *nonswitch* e-semantics (something like Teleological and Optimal e-semantics, respectively) with *switch* t-semantics. They are thereby committed to No/Yes views, which implausibly entail that Traveller makes mistakes about her current experiences, for example, judging at t_3 that the tomato looks green (because her judgment contents have changed), when in fact it looks red (because her visual contents have not).²⁴

To avoid this implausible commitment, Macpherson thinks such theorists must do one of two things: either “look elsewhere [than representationalism] for a theory of phenomenal character or reconsider their commitment to externalism” (2005, 150). I think Macpherson is wrong to assume that renouncing representationalism would help. Even if visual color content did not determine character, it would be implausible for Yes/No theorists to concede that Traveller does not know at t_3 what color the tomato looks to her, what color her experience represents it to be, just as it is implausible that Sally does not know on Twin Earth what liquid she is thinking about. But in that case, it can begin to look as though t-externalism has to go.

But it does not. All that has to go is a *certain* t-externalism, and it only in certain cases. Notice that the challenge of achieving a No/No view or a tight Yes/Yes view looks difficult only given the idea that the determinants of the color contents of experiences, on the one hand, and thoughts, on the other, are mutually independent. What is needed, I shall now suggest, is a t-externalism that yokes those two levels of color content together.

9. Bottom-Up Externalism

The upshot of the preceding difficulties is not that Yes theorists or No theorists need to renounce externalism, but that they need to hold that the color contents of thoughts are determined by the color contents of experiences. Call this bottom-up t-semantics. This is an externalist account of color concepts, given an externalist account of visual color content. It can figure as part of a more general bottom-up approach (as I shall call it) to other thought contents. Not only is this a highly plausible account of color concepts, but No theorists can use it to secure a No/No view, and Yes theorists to secure a tight Yes/Yes view.

Macpherson acknowledges this bottom-up route to a No/No view. But she describes it as “limited” and dialectically “odd” and thinks it would require No theorists to “reject, in a potentially radical way, the traditional externalist conception of the propositional attitudes” deriving from Putnam and Burge (Macpherson 2005, 148–50). If Macpherson is right, Yes theorists too face hard choices, given my claim that they need bottom-up t-semantics to generate a sufficiently tight Yes/Yes view. Fortunately, however, the dialectical position is more hospitable than Macpherson makes it appear, both for Yes theorists and for No theorists.

For one thing, we need not choose between applying a bottom-up approach to all concepts or none. The approach only can be applied to concepts that specify perceptual contents, and below I make room for the idea that it might be applied only to a subset of those.²⁵ What we admittedly cannot do is apply to a given concept both a bottom-up approach and an account that, like Burge’s or Putnam’s, grants a determinative role to socio-linguistic factors, unless these figure in the e-semantics. But, compatibly with acknowledging the bearing of such factors on other concepts, I think that in the case of color concepts there are strong reasons for applying the bottom-up approach—and not only that it helps with switch cases. These reasons have weighed with others. For the approach is not as novel as Macpherson supposes (2005, 150), being found in the writings of Brewer (1999), John Campbell (1993), Evans (1982), and John McDowell (1994).²⁶ In short, then, it would be helpful, independently plausible, and not unduly costly in terms of their other

externalist commitments for both Yes theorists and No theorists to adopt bottom-up t-semantics. Let me conclude, then, by sketching a little further the view's shape and motivations. A more detailed examination must await another occasion.

A natural way to develop a bottom-up approach is to claim that possession of the target concept consists in possession of a recognitional capacity. The referent of the concept is whatever property the capacity is a capacity to recognize. This in turn is determined by the content of the experiences grounding the capacity. Historically, many have held that perceptual experiences not only justify thoughts but fix their contents. But a bottom-up approach to a given concept, however, distinctively takes the content-determining feature of the grounding experiences to be *their* content, in particular, their representation of the very property the target concept is a concept of. Contrast Peacocke's view (1983, 37) that possessing the concept *red* involves being disposed to deploy it on the basis of experiences enjoying a certain quale; and Brown's plausible idea (1998) that some thinkers' possession of the concept *water* involves their having a water-recognizing capacity grounded in perception of water's superficial properties.

Beyond its making sense of switch cases, what are the motivations for taking a bottom-up approach to color concepts? One idea is that our concepts can refer to objective particulars and properties only if our perceptual experiences constrain our thinking, in the sense of determining the content of at least some of our concepts by way of being the canonical grounds for their deployment. This suggests we should adopt a bottom-up approach where we can, on pain of leaving our repertoire of concepts lacking empirical content.²⁷ A related idea, anticipated earlier, is that experience enables us to meet the "know-which" requirement on thinking, that is, the requirement that, with respect to at least some of our concepts, we can possess them only if we know which items are their referents.²⁸ In the case of particulars, the idea is that perceptual experience enables singular demonstrative thinking by enabling one to *locate* the objects of one's thought; in the case of certain basic properties, it is that perceptual experience enables one to discriminate and recognize instances of the properties in question. Arguably what is needed to meet the requirement in these basic cases is that experience determines reference in the sense that a difference in the reference of two concepts entails a difference in the reference of perceptual contents on the basis of which the concepts are deployed. This is precisely what the bottom-up t-semantics provides in the case of color concepts.

Given these motivations, it is plausible to restrict the bottom-up approach to concepts of what I call *appearance properties*, that is, properties whose perceptual representation *grounds* their recognition. Arguably, redness is an appearance

property and being water is not, for—to gesture at the argument—there could not be fake red things in the sense that there can be fake water, that is, nonwater we might mistake as water even in optimal perceptual circumstances on the basis of veridically perceiving its other properties.

Even if being water is not an appearance property, notice, one might yet allow that the concept *water* specifies visual contents. If one did allow this, and applied the bottom-up approach to *red* but not *water*, the result would be a bottom-up view of a *subset* of concepts specifying perceptual contents, as anticipated above. Would this account risk either a Yes/No or No/Yes view of switch cases involving water contents, hence allowing mistaken judgments about whether one’s own experiences are currently water-representing? Not if one adopted a *top-down* approach to water contents, on which one’s experiences are water-representing rather than, say, *twater*-representing in virtue of its being the case that the term one would use to report the experience—for example, “water”—expresses on one’s lips *water* rather than *twater*.

None of this means that sociolinguistic factors do not bear on the possession of noncolor concepts. We can recognize the relevance of such factors while acknowledging, to put it telegraphically, that a thinker’s conceptual repertoire cannot be put in contact with the world *entirely* courtesy of deference to a linguistic community, and that color concepts are points at which perceptual experience rather than the community plays the determinative role. This chimes with our intuition that the congenitally blind cannot acquire by deference the color concepts the sighted apply on the basis of their visual experiences.²⁹

These remarks about bottom-up t-semantics are only a sketch, of course, merely brushing the surface of deep and thorny issues. But even with this amount of flesh on its bones, the bottom-up approach to color concepts looks plausible, and a promising way for No theorists to secure a No/No view and Yes theorists a sufficiently tight Yes/Yes view. So I conclude that, as far as the objections considered in this paper go, both externalist No views and Yes views remain viable. This is good news for the ORV account of color. ORV theorists can embrace externalism, as required, and let their choice of a specific e-semantics determine their answer to the visual change question.³⁰

Notes

¹ Arguably our experiences represent not determinable colors, but determinate shades; however, nothing in what follows hangs on this issue.

² For the sake of argument, I grant ORV’s opponent that such inverting lenses—or, more plausibly, some more invasive mechanisms of inversion—are empirically possible, or at least that their empirical impossibility would not undermine the philosophical point.

³ On why it would be a pseudo-question, see Peacocke 1988.

⁴ The cabbage point requires that experiences of nausea do not represent the cabbage, not that they are representationally blank.

⁵ This stipulation—which, at least for the sake of argument, I grant ORV’s opponent—is more credible if Abe is ignorant of the lenses, perhaps because the objects he sees are unfamiliar or because he lacks standing color beliefs.

⁶ Campbell (1993) presents the symmetry problem differently but draws the same moral.

⁷ Tye (2000) regards the reference to optimal circumstances as operating like a *ceteris paribus* clause in an empirical generalization.

⁸ Actually, Block also takes into account the experiences’ “output,” i.e., thoughts and actions (1990, 58). It will simplify matters, however, and not affect the arguments that follow, to think of him as endorsing Typical e-semantics.

⁹ Relying on context, (i) “That” occurs in these judgments as a memory-based demonstrative. (ii) In the remainder of the paper, I use symbols from this table, e.g., “A₂,” to refer either to the sentences listed or to the judgments they express. The context disambiguates.

¹⁰ See Block 1994, 518; 1996, 42. If philosophically literate, Traveller might also utter D₂ (Block 1996, 42, 45). It is to ensure Traveller remains internally identical to Untravelled, hence has these dispositions, that Block moves her to Inverted Earth. The same effect could be achieved by supposing Traveller suffers amnesia, or that she and her twin lack beliefs about the colors of familiar types of object, or that she and her twin have around them no such types. The point of Inverted Earth is *not* to ensure Traveller remains ignorant of the lenses, which ignorance is later dropped in Block’s 1996 version of the thought experiment.

¹¹ Block himself appears to miss the argument to character independence (1996, 46, n. 8; 1990, 61). But it follows from his argument for content independence.

¹² (a) I take it Peacocke, by contrast, would say that the first and second experiences differ not only in color content, but color qualia. One phenomenal difference Peacocke clearly takes to be a difference of content *and not qualia* is that between experiences of a wire cube before and after a Necker-cube style aspect switch (1983, 16–17). Notice that the intentional difference in that case is in spatial content, however, not color content. (b) The notion of *color* character may need to be elaborated further to figure in the mixed view. In any case, the notion is dispensable; the mixed view and representationalist views can be formulated without it.

¹³ The same applies to “red-feeling” if qualia alone constitute character.

¹⁴ See also Lycan 1996, 110–11, 122–25.

¹⁵ On Evans’s procedure not yielding knowledge of qualia, see my 1999 (35, 62) and Byrne 2001.

¹⁶ It may well be more felicitous to speak of one’s having a memory of the tomato and its color, from a certain perspective, rather than one’s having a memory of the experience. But nothing here hangs on the issue.

¹⁷ Their point is part of a larger case against physicalist versions of ORV, a case weakened, incidentally, by their neglect of No views.

¹⁸ Boghossian and Velleman use “appear” but mean what I mean by

“look.” They use “looks red” as I am using “red-feeling” (1991, 69).

¹⁹ Goldman’s final formulation of his requirement emphasizes reliability, not discrimination, but he does not envisage the two coming apart.

²⁰ It is a nice question whether this threatens noncomparate knowledge also in the case of those who are not switch victims. If Sally cannot introspectively discriminate water and twater thoughts, nonvictims could not either. But given Twin Earth does not actually exist, arguably a twater thought is not a *relevant* non-water-thought for nonvictims.

²¹ Brown (2004) exploits the notion of relevance to overcome the problem not only for nonvictims but for Sally too. See also note 20.

²² Boghossian and Velleman 1991 is a rich article with strands not discussed here, e.g., their arguments concerning our knowledge of color resemblances (91–97) and the possibility of red things failing to constitute a natural kind (91).

²³ I say “perhaps” for two reasons. First, the risk of a No/Yes view in this case exists only if Untravelled is not alerted to the replacement. (If she is, she can either resolve to continue using “red” in her original community’s way, or intentionally adopt the newcomers’ language, in which case she will at t_3 say, correctly, “The poppy looks green[red].”) But if we ensure her ignorance of the replacement by supposing that she and the newcomers lack standing color beliefs, and never have an opportunity to compare perceptual color judgments about the same object, it might be doubted whether the newcomers become in the relevant sense her new linguistic community with respect to color terms. Second, if one thinks a change in allegiance must be intentional, then Untravelled’s ignorance of the replacement excludes a change in allegiance. Notice that Block could not take this line, however, since he allows an unintentional change in allegiance in the original Inverted Earth case.

²⁴ Lycan fails to rule out a No/Yes view in his 1996 (134).

²⁵ In footnotes, Macpherson acknowledges the necessary condition but appears to rule out further restrictions (2005, 147, n. 37; 148, n. 40).

²⁶ This is not to say that they would agree with every aspect of the following sketch.

²⁷ See Brewer 1999 and McDowell 1994.

²⁸ See Brewer 1999, Campbell 1993, and Evans 1982.

²⁹ Is the bottom-up approach available to disbelievers in nonconceptual content (call them conceptualists)? Arguably so. Bottom-up semantics says that possessing the concept *red* requires sometimes undergoing experiences representing things as red. Some suspect a vicious circle is completed if conceptualists add that having such experiences in turn requires possession of *red* (Peacocke 1992; contrast McDowell 1994). But a conceptualist might point out that the underlined occurrence of “red” does not refer to the concept *red*, nor to its being possessed. It refers to a color and only *displays* the mode of presentation under which things must sometimes be visually presented to one if one is to possess *red*.

³⁰ For discussion of this paper and earlier material on which it is based, I am very grateful to Carol Bain, George Bain, Bill Brewer, Jessica Brown, Bill Child, Rose Drew, Fiona Macpherson, Philip Percival, Adam Rieger, Alan Weir, and audiences at the Universities of Bristol and Kent.

David Bain

Bibliography

- Bain, D. 1999. *Sensation and representation: A study of intentionalist accounts of the bodily sensations*. DPhil Thesis. Bodleian Library, University of Oxford.
- Block, N. 1998. Is experiencing just representing? *Philosophy and Phenomenological Research* 58:663–70.
- . 1996. Mental paint and mental latex. In “Perception” *Philosophical Issues* 7:19–49.
- . 1994. Qualia. In *Companion to the philosophy of mind*, ed. S. Guttenplan. Oxford: Blackwell.
- . 1990. Inverted earth. In “Action theory and philosophy of mind,” ed. J. E. Tomberlin. *Philosophical Perspectives* 4:53–79.
- Boghossian, P., and D. Velleman. 1991. Physicalist theories of color. *Philosophical Review* 100:67–106.
- Brewer, B. 2001. Review of Michael Tye’s *Consciousness, colour, and content*. *Mind* 110:869–74.
- . 1999. *Perception and reason*. Oxford: Clarendon Press.
- Brown, J. 2004. *Anti-individualism and knowledge*. Cambridge, MA: MIT Press.
- . 1998. Recognitional capacities and natural kind terms. *Mind* 107:275–303.
- Burge, T. 1979. Individualism and the mental. *Midwest Studies in Philosophy* 4:73–122.
- Byrne, A. 2001. Intentionalism defended. *Philosophical Review* 110:199–240.
- Campbell, J. 1993. A simple view of colour. In *Reality, representation, and projection*, ed. J. Haldane and C. Wright. Oxford: Oxford University Press.
- Dretske, F. 1995. *Naturalizing the mind*. Cambridge, MA: MIT Press.
- Evans, G. 1982. *Varieties of reference*. Oxford: Oxford University Press.
- Goldman, A. 1976. A causal theory of knowing. *Journal of Philosophy* 73:771–91.
- Lycan, W. G. 1996. *Consciousness and experience*. Cambridge, MA: MIT Press.
- Macpherson, F. 2005. Colour inversion problems for representationalism. *Philosophy and Phenomenological Research* 70:127–52.
- McDowell, J. 1994. *Mind and world*. Cambridge, MA: Harvard University Press.
- Millikan, R. G. 1984. *Language, thought, and other biological categories*. Cambridge, MA: Cambridge University Press.
- Peacocke, C. 1992. *A study of concepts*. Cambridge, MA: MIT Press.
- . 1988. The limits of intelligibility: A post-verificationist proposal. *Philosophical Review* 97:463–96.
- . 1983. *Sense and content*. Oxford: Oxford University Press.
- Putnam, H. 1975. The meaning of ‘meaning’. *Minnesota Studies in the Philosophy of Science* 7:131–93.
- Tye, M. 2000. *Consciousness, color, and content*. Cambridge, MA: MIT Press.