

It's Not Easy Being Green: Hardin and Color Relationalism

Jonathan Cohen*

It's not that easy being green;
Having to spend each day the color of the leaves.
When I think it could be nicer being red, or yellow or gold...
or something much more colorful like that.

— Kermit the Frog

C. .L. Hardin is that rarest of things — a philosopher who has changed the world. Prior to the publication of *Color for Philosophers: Unweaving the Rainbow* (Hardin, 1988), philosophical work on color had been conducted in roughly the same terms in which it was carried out by the famous moderns — Galileo, Boyle, Locke, et. al. But once Hardin pointed out that a vast field of empirical research had developed since the modern period, and showed convincingly that these developments impose serious constraints on ontological and epistemological disputes about color, the philosophical landscape was forever changed. Subsequently, philosophical work on color has increased dramatically in both sophistication and interest (as shown by its growth in recent years), and we have Hardin to thank for these salutary developments.

But Hardin hasn't contented himself with reframing traditional philosophical issues about color in a way that is sensitive to relevant empirical constraints. In addition, he has been a staunch defender of color eliminativism — the view that there are no colors, *qua* properties of tables, chairs, and other mind-external objects, and a vociferous critic of several varieties of realism about color that have been defended by others (e.g. Hardin, 2003, 2004). Among the realist views that Hardin has targeted are the so-called color physicalism of Hilbert (1987); Byrne and Hilbert (1997a, 2003); and Tye (2000),¹ and, inconveniently for me, the relationalist views of Cohen (2003a, 2004, 2003b); McLaughlin (2003); and Jakab and McLaughlin (2003).

*Department of Philosophy, University of California, San Diego, 9500 Gilman Drive, La Jolla, CA 92093-0119, joncohen@aardvark.ucsd.edu

¹'Physicalism' strikes me as an inapt label for the view of Hilbert et. al., insofar as (i) alternative accounts also allow that colors are physical, (ii) the view presently under consideration comes without any substantive characterization of the physical, and (iii) the issue about what counts as physical is, in any case, orthogonal to the main lines of contrast concerning the nature of color. That said, I'll adhere to the label preferred by the defenders of the view for present purposes.

My main purpose in this paper is to defend relationalism against Hardin's most recent criticisms of it (Hardin, 2004). However, to motivate this discussion, it will be helpful first to consider Hardin's recent objections against color physicalism, and why existing physicalist defenses against such objections are unsatisfactory (§1). This will put us in a position to see some of the virtues of color relationalism (§2), and then to evaluate Hardin's complaints against it (§3). I'll argue that relationalism survives Hardin's criticisms, and that it is ultimately preferable to both color physicalism and Hardin's own eliminativist view.

1 Hardin on Color Physicalism

Many philosophers inclined toward realism about color have proposed that (surface) colors are some sorts of mind-independent properties of object surfaces. Recent color physicalists have defended the more specific view that (surface) colors are classes of surface spectral reflectance distributions.²

Hardin and others have pointed out that this form of color physicalism faces significant hurdles stemming from the remarkably widespread interpersonal and intrapersonal (and interspecies) perceptual variation with respect to color: a given stimulus produces a variety of effects in different perceptual systems, and produces a variety of effects on a single perceptual system when viewed under different perceptual circumstances. On standard assumptions that are common ground in the current dispute, each of these different effects is a representation of the color of the stimulus.³ But if colors are mind-independent and circumstance-independent properties of surfaces, as are spectral reflectance distributions (or classes thereof), then physicalists are committed to saying that at most one of these varying effects represents the color of the stimulus veridically. However, the objection goes, it is hard to see that anything could (metaphysically) make it the case that one of the variants is veridical at the expense of the other: it seems that any considerations that could be brought forward in support of the

²A surface will reflect some percentage of the light of wavelength λ that falls on it. If we collect the percentages of reflected to incident light for each visible wavelength, we will get a function (from visible wavelengths to numbers in the interval $[0,1]$) that characterizes the disposition of a surface to affect light in the visible range in a certain way. This function is the surface spectral reflectance distribution.

Physicalists typically prefer to identify colors with classes of such functions (rather than with the functions themselves) because of the phenomenon of metamerism: under any given illumination, an infinite number of surfaces (distinct in their surface spectral reflectance distributions) will be visually indistinguishable for a given observer. What this suggests is that identifying colors with reflectance functions yields an excessively fine-grained individuation of the colors. The move to identify colors with classes of reflectance functions is intended to get around this problem. (An exception to this generalization is Churchland (2007), who identifies colors with reflectances, and is prepared to live with the resultant extremely fine-grained individuation of colors.)

³These standard assumptions are rejected by, e.g., Smith (2002); Travis (2004), who deny that vision and color experience represent the stimulus and its properties in the relevant way. Since, as noted in the main text, these assumptions are not disputed by the main parties to the current dispute, I'll continue to ignore this issue in what follows.

veridicality of one of the variants could be matched by considerations of equal force in favor of some other variant.

As Hardin (2004) hammers home, the sort of variation of color vision at issue is no mere imagined possibility. For one thing, there is overwhelming and unambiguous evidence of actual variation of color vision in normal human subjects (and even more when we look at non-human visual systems), holding all viewing conditions (e.g., illumination, background, viewing distance, angular subtense of visual field, etc.) fixed. One among many clear instances of this sort of interpersonal variation, discussed at length by Hardin (2004), is the variation in the spectral wavelength (alternatively, in the Munsell chip) selected by subjects as unique green (i.e., as looking greenish without looking at all bluish or at all yellowish). When two normal trichromatic observers view chip *C* under identical perceptual conditions, *C* looks unique green to one of them but bluish green (hence not unique green) to the other. For that matter, as described by Cohen (2004) and Hardin (1988), there is equally overwhelming and unambiguous evidence of intrapersonal variation of color vision as a function of many parameters of the viewing condition. Consequently, there are cases of intrapersonal variation that mirror the interpersonal one just discussed. For example, a chip *C* can look greenish to you when presented against one background and bluish when presented to you against another background. If color physicalism is true, then in this intrapersonal case as well, it must be that at most one of the competing representations veridically represents *C*'s color. But, once again, it is extremely hard to imagine what could (metaphysically) make it the case that one of the representational variants is veridical at the expense of the other.⁴

Physicalists have sometimes suggested in response to this concern (e.g. Byrne and Hilbert, 2004, 2003; Tye, 2000) that the difficulty here is merely epistemic — that while facts about variation might prevent us from *knowing* which of the competing variants is veridical at the expense of the others, there is nonetheless a (possibly unknown) fact of this matter, just as color physicalism requires. But, as far as I can see there is literally no reason to believe that that is true. On the contrary, there is substantial (but defeasible) reason to believe that it is false — namely, the failure of several hundred years of systematic efforts directed at uncovering such facts of the matter establishes a presumptive case against their existence. As such, the physicalist's insistence that there is an epistemically unavailable fact of the matter strikes me as a piece of unwarranted optimism.⁵⁶

⁴Significantly, the similar structure of the problem that arises in both interpersonal and intrapersonal cases of perceptual variation suggests that we should be seeking a single solution that applies to both sorts of cases.

⁵Cohen (2007) elaborates this response, and provides some additional reasons for doubting that there is some unknown fact of the matter that would sustain the physicalist's commitment.

⁶Hardin (2004, 35) offers a further, if less developed, complaint against the color physicalist's insistence that there is an unknown fact of the matter about which perceptual variant veridically represents the color of a given stimulus. Namely, he suggests, if the color physicalist is driven to posit unknown (or even unknowable) facts about colors, it restores just the sort of veil of perception between us and the world that, one might have hoped, a physicalist ontology held the promise of sweeping away, and thereby undercuts one of the initial attractions of the view. I'll put this point aside, since I am not convinced that such epistemological

2 Color Relationalism

As we saw in §1, color physicalists are committed to claiming that at most one among conflicting perceptual variants veridically represents the color of a single stimulus. The difficulty with this commitment, I suggested, is that there is no reason for believing that any one of the variants is distinguished from the others, as the physicalist maintains. In the face of these difficulties, the relationalist suggests that we should avoid the trouble by refusing to choose between the variants. That is, she suggests, the way out of the trouble is to hold that the conflict between the variants is only apparent, insofar as the single stimulus is genuinely both unique green to observer S_1 and not unique green to S_2 . Likewise, since a single stimulus can look unique green to a single observer S under one viewing condition C_1 but fail to look unique green to S under another viewing condition C_2 , and since there seems not to be any fact of the matter that makes one of these two representations of the stimulus's color veridical at the expense of the other, the relationalist will refuse to choose between the two: instead, she'll insist, the stimulus is genuinely both unique green to S under C_1 , and not unique green to S under C_2 .

What the relationalist proposes, then, is that colors are not (as the physicalist maintains) subject- and condition-independent properties of their bearers, but relational properties that are constituted in terms of relations to subjects and viewing conditions. Since, on this view, colors are relational properties, the perceptual variants that we initially characterized as conflicting are in fact not conflicting. This view does justice to the facts about perceptual variation, and it does so without requiring either unmotivated choices between variants or unjustified optimism that there is some unknown (or unknowable) fact that could motivate such a choice. This, it seems to me, is an important virtue of the view, and one that makes it worth taking seriously for those who aspire to realism about color.

That said, it is worth emphasizing that relationalism is not, by itself, a theory of the nature of color. It is a theory about what sorts of properties colors are — namely, that they are relational properties; but it does not say which properties of that sort — which relational properties, in particular — colors are. In fact, there are several species of relationalism; these include the dispositionalist view of McGinn (1983); Peacocke (1984), and Johnston (1992), the so-called “enactive” view of Thompson *et al.* (1992); Thompson (1995), the “selectionist” view of Matthen (2005), the (distinct) “color functionalist” positions of McLaughlin (2003); Cohen (2003a), and others. The question of whether colors are relational puts substantive constraints on what counts as an adequate color ontology, but answering that question will leave plenty of room for disagreement about what colors are. Consequently, relationalism might aptly be regarded as a family of accounts of color ontology — a family whose members share some commitments but not others, rather than as an account in its own right. (I'll return to this theme in §3.2.2.)

motivations are as essential in motivating color physicalism as Hardin maintains, and since I think physicalism faces other serious problems (as discussed in the main text).

3 Hardin on Relationalism

As we have seen, relationalist views avoid at least one important class of difficulties that plague color physicalism. However, Hardin presses further worries against relationalism. First, he urges that the view doesn't deliver the kind of objectivity that a color realist should want, and so doesn't amount to an acceptable form of color realism after all. Second, he worries that color relationalism is unacceptably liberal about the attribution of colors. I'll take these worries in turn (§§3.1–3.2), and then turn to a conjecture about what might ultimately underlie Hardin's other reasons for being worried about relationalism (§3.3).

3.1 Objectivity Lost?

Hardin's first concern about color relationalism is that the view makes colors insufficiently objective — no more objective, in any case, than properties like *being beautiful* and *being ugly*. If so, he suggests, the kind of objectivity relationalism can bestow to color properties is too attenuated for a realism worthy of the name:

To me, this woman's face is the very Form of Beauty incarnate. To Jonathan, it is a face only a blind mother could love. Beauty is in the eye of the beholder, you say? No, it is an objective property of the woman's face, for I need only to relativize it to Larry's gaze, if only at time T . Ugliness is also an objective property of her face, provided of course we understand it as being relative to Jonathan's eye at time T . And so this same woman has the possibility of being all things to all men (Hardin, 2004, 36).

Is Hardin right that relationalism threatens to make *being beautiful* an objective property also? Obviously, the answer to that question depends rather a lot on what one means by 'objective', and unfortunately there are a variety of candidate understandings with a historically salient claim on the word.

It does seem that there are several relevant understandings of 'objectivity' on which *being green* and *being beautiful* (as understood by the relationalist) turn out to be objective. For example, suppose a property P is objective just in case individuals in the world genuinely exemplify P ; for the relationalist about *being green* and *being beautiful*, both properties are objective in this sense, insofar as things genuinely are green/beautiful (viz., by virtue of standing in certain relations to subjects). Alternatively, suppose we count P as objective just in case P would continue to have its instances if (counterfactually) there were no perceiving subjects; here, again, a relationalist can hold that both properties are objective in the sense at issue in that the relations in virtue of which things count as green/beautiful can be spelled out in terms of how those things *would* (possibly counterfactually) strike subjects.

In addition, there are other historically salient understandings of 'objective' that underwrite the opposite verdict about both properties. For example, objectivity is sometimes understood in terms of interpersonal convergence; e.g.,

one might think of P as an objective property just in case undistracted and knowledgeable subjects agree in judging whether x is P for a wide range of x . Of course, one of the lessons of §1 is that color properties fail this criterion; and, of course, many (e.g., Hardin, in the quotation above) have thought the same is true of *being beautiful* (but see, e.g., Mothersill, 1984; Zangwill, 2001). Another criterion for property objectivity, built on the traditional subject-object contrast, might hold that property P is objective if and only if it is constituted independently of subjects. Of course, relationalism about P just amounts to denying that this criterion holds of P , so the relationalist about *being green* and *being beautiful* will hold that both properties are non-objective on the present understanding.

So there are several criteria for objectivity that agree in treating *being green* and *being beautiful* as objective properties, and there are several others that agree in treating them both as non-objective properties. On the other hand, it is not inevitable that the color relationalist will classify the two properties on the same side of the objective/non-objective distinction.

First, suppose that, as before, we accept relationalism about both *being green* and *being beautiful*, but that we understand property objectivity in terms of the subject-independence of its causal effects; in particular, suppose a property P counts as objective just in case the event of its exemplification by x has causal effects that are independent of the effects x has on subjects. As far as I can see, this further (and perfectly good) sense of property objectivity leaves it entirely open whether distinct relational properties (*being green* and *being beautiful*) are alike or different in respect of their objectivity, depending on the contingencies of causal pathways in the world. Moreover, and independently of this last point, it is important to recall that a color relationalist is not *per se* committed to a relationalist treatment of *being beautiful*; consequently, if she favors an alternative view about the latter, many of the criteria for property objectivity listed above may count color properties as objective and *being beautiful* as non-objective, or vice versa.

It seems, then, that, contrary to what Hardin suggests, the color relationalist is forced to conclude neither that *being green* is non-objective, nor that it is only as objective as *being beautiful*. But suppose (contrary to fact) that a relationalist were committed to regarding colors as no more objective than *being beautiful*. Why would that be an objectionable conclusion? Of course, there are some understandings of ‘realism’ that would have the effect of undermining the relationalist’s realist aspirations — namely, those that make criterial one of the senses of objectivity that are unavailable to the relationalist. On the other hand, this point itself seems a reason for being dissatisfied by those construals of ‘realism’, rather than taking it as an indictment of relationalism.⁷ After all, insofar as relationalism is precisely a theory of how colors are constituted, relationalists are committed to saying that colors are *bona fide* properties that are (sometimes) exemplified by *bona fide* instances — otherwise there would be no target for relationalism to be a theory about in the first place (cf. Sayre-

⁷I’ll return to this point in §3.3.

McCord, 1988; Sober, 1982; McDowell, 1985). So far as I can see, then, there is no reason that a relationalist need fear the conclusion that *being green* is only as objective as *being beautiful*.

3.2 A Chromatic Swarm?

Hardin has a second reason for being unhappy with color relationalism that turns on the plurality of color properties recognized by the view. In this connection he employs a venerable philosophical strategy — the argument by quotation. In this case, he quotes Meno’s explication of Gorgias’s definition of virtue:

First of all, if it is manly virtue you are after, it is easy to see that the virtue of a man consists in managing the city’s affairs capably, and so that he will help his friends and injure his foes while taking care to come to no harm himself. Or if you want a woman’s virtue, that is easily described. She must be a good housewife, careful with her stores and obedient to her husband. Then there is another virtue for a child, male or female, and another for an old man, free or slave as you like; and a great many more kinds of virtue, so that no one need be at a loss to say what it is. For every act and every time of life, with reference to each separate function, there is a virtue for each one of us, and similarly, I should say, a vice (Plato, *Meno*, 71E).

The argument finishes off by quoting Socrates’s sarcastic reaction to Meno’s disquisition: “How fortunate I am, Meno! I wanted one virtue and I find that you have a whole swarm of virtues to offer.”

Despite the exegetical perils attendant on the attempt to understand arguments by quotation (especially those quoting such vexed texts as the *Meno*), I think I can discern two possible anti-relationalist objections that Hardin might have in mind here. First, one might worry that relationalism is objectionable on strictly numerical grounds: the thought would be that relationalism sustains too many color attributions – that it results in what Hardin at one point terms “ungainly pluralities” (36). A second possible concern, which I take to be closer to what underlies Socrates’s objection to Gorgias’s account of virtue, is that the distinct colors recognized by the relationalist fail to be appropriately unified. Once again, I’ll take these concerns one at a time.

3.2.1 Ungainly Pluralities

It seems hard to deny that there is something *prima facie* counterintuitive in the relationalist’s claim that a given stimulus can exemplify (at the same place, at the same time) multiple colors, each relativized to different subjects and perceptual conditions. Most significantly, this relationalist claim seems to fly in the face of our ordinary conversational practice of attributing a single color to a given object at a given time: we say that the ripe lime is green, and leave it at that, and speak without hesitation about *the* color of the lime. There is, I

think, no serious question about these facts about the way we attribute colors. The only serious question is how damaging these facts are to color relationalism.

My reason for thinking that they are not so damaging after all is that data about the property attributions occurring in ordinary conversation are the joint result of (i) facts about the properties figuring in the attributions, and (ii) the communicative (and other) presuppositions we bring to the linguistic exchanges in which the attributions occur. I suggest that the apparently troublesome data about color attributions are to be explained in terms of factors of type (ii) rather than type (i), and so are not ultimately at odds with color relationalism *qua* account of the nature of color properties.⁸

That is, the reason we typically attribute in conversation only a single color to a particular ripe lime at time t is that, while the lime exemplifies many colors at t (it is green to S_1 in C_1 , red to S_2 in C_2 , orange to S_3 in C_3 , and so on), we only care about a very small finite minority of those colors — in many cases only one of them — in typical conversational settings. For one thing, in such settings we are generally unconcerned with the colors that the lime exemplifies in virtue of its relations to visual systems very much unlike our own (e.g., octopus visual systems, imagined Martian visual systems, and even human dichromatic and anomalous trichromatic visual systems), and so don't mention these colors that (according to the relationalist) the lime genuinely exemplifies. And although there will be non-trivial variation between the colors the lime has relative to visual systems we do care about (see §1) that is discernible through systematic testing in the psychophysics lab, this variation won't reveal itself in ordinary conversational settings outside the lab; consequently, ignoring this variation by the use of a more general color term — e.g., green *simpliciter* — will serve well enough for ordinary conversational purposes. Likewise, in such settings we are generally unconcerned with the colors that the lime exemplifies in virtue of relations to viewing circumstances substantially unlike those we typically encounter, or in which our interests in making color classifications are not in play (e.g., circumstances of extremely low illumination, or viewed through a chromatic filter, or where the stimulus subtends 0.01 degrees of visual angle, or against a background of diagonal achromatic lines and presented for 20 milliseconds), and so don't mention these colors that (according to the relationalist) the lime genuinely exemplifies. And although there will be non-trivial variation between the colors the lime has relative to visual circumstances we do care about (see §1) that is discernible through systematic testing in the psychophysics lab, this variation won't reveal itself in ordinary conversational settings outside the lab; consequently, ignoring this variation by the use of a more general color term — e.g., green *simpliciter* — will serve well enough for ordinary conversational purposes.

My proposal, then, is that the presuppositions we bring to ordinary conversational settings in which we make color attributions include general ways of filling in the parameters to which (according to relationalism) colors are always relativized. In effect, when we say that the lime is green, we mean that the lime

⁸Cohen (2004, §4) develops this idea in greater detail.

is green to visual systems pretty much like our own, under visual circumstances pretty much like those we typically encounter. These parameters go unmentioned in ordinary conversation not because they are not there (a conclusion that would be damning to color relationalism) but because the presuppositions in force in these settings allows us to leave them tacit (a conclusion that is consonant with color relationalism). Of course, if this strategy serves ordinary purposes, this leaves it open that there are some non-ordinary purposes (e.g., those in the psychophysics lab) in which the simplifications break down. This is, I suggest, exactly what we find: when systematic variation of the relevant parameters reveals the contribution they make to the color of a fixed stimulus, we make explicit the relationality of colors that had heretofore remained tacit — e.g., we say that the chip is unique green to *this* observer, but not *that* one, in *this* visual circumstance, but not *that* one.

By way of analogy, consider the weather. I take it that meteorological status is obviously constituted in terms of a relation to locations, which means that there really is a swarm of meteorological conditions: it is cold and raining in location L_1 , it is hot and dry in location L_2 , etc. But, so long as I don't investigate the systematic effects on weather of considering different locations, conversational presuppositions constrain the way we fill in the relevant relativizations; consequently, my neighbors and I can get by asking each other (not explicitly relativized) questions like "what is the weather today?". However, this strategy fails as soon as more wide-ranging investigation begins: when I talk to friends across town, I notice that what they say about the weather is very different from what I would say, and this effect increases as I vary the value of the locational parameter more dramatically. For example, Hardin lives in Syracuse and I live in San Diego, so he'll always say that it is raining and overcast, and I'll always say that it is sunny and 72°F. When we notice that this occurs, we make explicit the relationality of the meteorological situation that had heretofore remained tacit — e.g., we say that the weather is sunny and 72°F in San Diego, but not in Syracuse.

It seems, then, that dramatic meteorological variation can often be ruled out effectively by conversational presuppositions (e.g., those shared by my neighbors and me) that restrict attention to a small range of locations. This does not mean that there is no meteorological variation left in play even when such presuppositions are in force — there really can be meteorological variation over small changes in location. But it does mean that the variation in question is much less dramatic than it would be were they not in force, and that the residual variation is small enough that it is unlikely to manifest itself to casual observation (as opposed to detailed meteorological investigation).

So too, I suggest, dramatic perceptual variation with respect to color can often be ruled out by conversational presuppositions (e.g., those between normal human observers in the sorts of viewing conditions we typically encounter) that restrict attention to a small range of visual systems and viewing circumstances. This does not mean that there is no perceptual variation with respect to color left in play even when such presuppositions are in force — there really can be perceptual variation with respect to color between normal human visual systems

and in the sorts of viewing conditions we typically encounter. But it does mean that the variation in question is much less dramatic than it would be were they not in force, and that the residual variation is small enough that it is unlikely to manifest itself to causal observation (as opposed to detailed psychophysical investigation).

For these reasons, I think color relationalism can reasonably be reconciled with our ordinary practices of color attribution, and that this defuses the numerical worry we have been considering.

3.2.2 Unity of the Colors

Let us return, then, to the other worry discernible in the passage Hardin quotes, which concerns the unity of the colors. As I read the text, the heart of Socrates's objection to Gorgias's view of virtue is not that it results in a mere plurality of virtues, but that there is no principled criterion for inclusion in the plurality. As Socrates puts the point a few lines down from the quoted passage, "Even if they are many and various, yet at least they all have some common character which makes them virtues. That is what ought to be kept in view by anyone who answers the question, What is virtue?" (72d). If this is the right way to read Socrates's concern about virtue, then the analogous worry for the color relationalist is not merely that the view proliferates colors, but that there should be something — and something that it is the central task of a theory of color to elucidate — shared by the several colors recognized by the theory in virtue of which they do (and other properties do not) count as colors.

I am sympathetic to Socrates's objection to Gorgias (at least as I have construed it); I think it is fair to ask that a theory of virtue should amount to more than a motley list — that it should provide a principled and projectible criterion that unites those things it recognizes as virtues and distinguishes them from those it does not. Moreover, I share Socrates's concern that Gorgias's account, as related by Meno, does not meet this desideratum. If this is right, it seems only fair to insist that an adequate account of color should, likewise, provide a principled and projectible criterion that unites those things it recognizes as colors and distinguishes them from those it does not. I want to argue that the color relationalist is, unlike Gorgias, in a position to meet this desideratum.

Unfortunately, saying why requires some caution, insofar as the relationalist's answer will depend on commitments that go beyond the strict bounds of relationalism — namely, on details of the specific form of relationalism that she holds. Recall that, as emphasized in §2, relationalism is not by itself an account of the nature of color. Relationalism has it that colors are relational properties constituted in terms of relations to subjects and viewing conditions; however, saying this leaves open the crucial question of exactly which relations to subjects and viewing conditions are relevant — and the latter question is one on which different relationalists hold sharply divergent views (along a number of different dimensions). What I want to insist on for present purposes is that (i) each relationalist has an answer to this question, and (ii) each answer to this question will provide a principled and projectible criterion that unites those

things that the theory recognizes as colors and distinguishes them from those things it does not. It follows from these two points that each relationalist will have an answer to the challenge under consideration, even if relationalism *per se* is agnostic between these answers.

To make this so-far schematic point clearer, I'll illustrate by adverting to the example of my own preferred version of relationalism — the color functionalist account I have defended elsewhere (Cohen, 2003a, esp. §1.3). Suppose, therefore, that you are a functionalist about color: you hold that x is green for subject S in condition C just in case (and by virtue of the fact that) x has the (multiply realizable) property that confers on its bearers the disposition to look green to S in C , that x is red for subject S in condition C just in case (and by virtue of the fact that) x has the (multiply realizable) property that confers on its bearers the disposition to look red to S in C , and so on for the other colors. Then you will say that *being green for S in C* and *being red for S in C* are united in being colors because of this fact: their instances count as instances by virtue of exemplifying a (multiply realizable) property that confers on its bearers the disposition to look a certain way to certain subjects in certain circumstances. But this still isn't sufficient for answering Socrates's worry about unity, because we still owe a story about what it is for something to look one of those certain ways to certain subjects in certain conditions — we need an account of *looks green to S in C* , *looks red to S in C* , and so on. To answer this question is to go beyond relationalism once again; indeed, it is to go beyond functionalism *qua* theory of the nature of color, and to defend a theory about the nature of color experience. As before, I take it that a number of options are open at this point, and will only provide one way of fleshing out what is needed for the sake of illustration. Say that x looks green to S in C just in case by visually attending to x in C , S is appropriately caused (in C) to have an experience of green; and say that *experience of green* is a type of state of the opponent-process system (viz., a neurocomputationally individuated state type) (*mutatis mutandis* for *looks red to S in C* and its ilk).

On the functionalist flavor of relationalism, then, the unity of colors comes from the unity of the set of color experiences; and (on the neurocomputational identity-theoretic strategy sketched above) the latter gets its unity — that is, the set is distinguished from other things in a principled and projectible way — by the type-individuations made available by a well-known neurocomputational theory (viz., opponent-process theory; Hurvich and Jameson (1957); Boynton (1979); Hurvich (1981); Hubel (1988)). We are therefore in a position to meet Socrates's challenge, for we have a principled and projectible criterion for being a color experience, and (derivatively) we have a principled and projectible criterion for being a color.

That said, I want to emphasize that my point here is more general than the particular relationalist-friendly proposals sketched above, and so should not be taken to stand or fall with those particular proposals. Relationalists are a varied lot. They agree that colors are constituted in terms of relations to perceivers and viewing conditions, but they disagree amongst themselves about *which* relations to perceivers and viewing conditions constitute colors. Still, every relationalist

(or, any rate, every relationalist who has worked out a full account of the nature of color) is committed to some or other answer to this question. But since any answer to the question will amount to a principled and projectible criterion for being a color, this means that every relationalist (who has worked out a full account of the nature of color) can provide for the sort of unity that Socrates complains is lacking from Gorgias's account of the virtues. Even if there is no single response to Socrates that would be acceptable to every relationalist, no relationalist (who has worked out a full account of the nature of color) will be without an answer to Socrates.

3.3 Relationalism and Reality

Though I have addressed the complaints about relationalism that Hardin (2004) made explicit, he has revealed in comments on an earlier draft of this paper that he has a more fundamental worry that underlies the other criticisms discussed, and that I think warrants discussion. Namely, Hardin worries that color relationalism is an insufficiently robust form of realism about color, and so would be more forthrightly cast as a species of irrealism/eliminativism.

Unsurprisingly given his views, Hardin endorses the relationalist's strategy of explaining what unites the instances of colors ultimately in terms of facts about perceiving subjects rather than (exclusively) in terms of perceived objects. Similarly, he commends the relationalist for recognizing the need to appeal to perceiving subjects in explaining central (and arguably essential) properties of colors such as the similarity relations that obtain between them and their unique-binary structure. From his point of view, however, accepting that colors are constituted in subject-involving ways (and thereby making available these explanatory strategies that he agrees are needed) is tantamount to giving up on realism about color. If colors and their most important features are explained in terms of what is in heads rather than what is in objects, this is for Hardin reason to regard our ascription of colors to objects as a false projection onto those objects of what is not (or is not wholly) there. Hardin urges this point in the following characteristically eloquent passage:

The color structure of external physical objects is, on the relationalist view, parasitic on the color structure that manifests itself in perception. If external objects were the bearers of color, one would suppose that it is they who would bequeath color structure to perception rather than the other way around. Indeed, if that were so, I might be prepared to gesture outwardly and say, "Lo, the colors!" But it is not so, and thus my gesture must be inward (Hardin, 2007, 4).

But if relationalism is, in fact, a closet irrealism, then Hardin's advice to relationalists is to join him in embracing that irrealism with a clear conscience. Thus, Hardin extends an olive branch: "Come home, Jonathan. All is forgiven" (Hardin, 2007, 5).

Although I am grateful for the forgiveness Hardin offers, and flattered to be invited to join his ranks, I must demur. For I continue to believe that there is an important sense in which color relationalism is a form of realism (and in which Hardin's view is not). I am, of course, happy to give up the word 'realism'. The more substantive and important issue is that, according to color relationalism (and not according to Hardin's eliminativism) colors *qua* properties that objects bear are among the true inventory of what there is. Indeed, one way to think about the lesson of relationalism about colors (or *being beautiful*, *being humorous*, and so on) is that, while what there is contains some elements constituted independently of subjects, there are other and equally good elements of what there is that are constituted in a way that depends essentially on subjects. In the particular case at hand, color relationalism teaches us that, contrary to what Hardin supposes in the quoted passage, external objects can be bearers of colors even if the unity and structure of colors is bequeathed to them by (subject-involving) perception.

Hardin's invitation depends on jettisoning from our ontology whatever is constituted in a way that depends on subjects. Therefore, it turns on what will seem to the relationalist an inappropriately restricted accounting of what there is. For this reason, I cannot accept his gracious invitation in good faith.

4 Conclusion

I think Hardin is right to reject color physicalism. That view makes commitments that we have empirically motivated reasons for doubting, and can react to these doubts only by an unwarranted optimism (§1). On the other hand, I think Hardin is too quick to reject relationalism about color. Relationalism has the virtue that it avoids the unmotivated choices and unwarranted optimism of color physicalism, and does justice to the full range of observed perceptual variation with respect to color (§2). Moreover, I have argued (§3) that relationalists can answer Hardin's concerns about the objectivity of the view, its commitment to ungainly pluralities, and its lack of unity, while continuing to recognize colors as *bona fide* parts of what there is. While relationalism might involve a revision of some of our naïve views about color, I propose that the view allows us to hold onto more of what we take to be central, while accommodating the empirical facts, than competing alternatives. In particular, it is a much less radical revision than Hardin's own eliminativist view that does away with colors *qua* properties of extra-mental entities entirely. As a famous philosopher has pointed out, it's not easy *being green*; however, the viability of relationalism shows that it is not impossible.⁹

⁹I am grateful to Larry Hardin, John Jacobson, Dom Lopes, Mohan Matthen, and Dana Nelkin for helpful comments on earlier drafts of this paper, and to an audience at the 2007 Society for Philosophy and Psychology.

References

- Boynton, R. M. (1979). *Human Color Vision*. Holt, Rinehart and Winston, New York.
- Byrne, A. and Hilbert, D. R. (1997a). Colors and reflectances. In A. Byrne and D. R. Hilbert, editors, *Readings on Color, Volume 1: The Philosophy of Color*, pages 263–288. MIT Press, Cambridge, Massachusetts.
- Byrne, A. and Hilbert, D. R., editors (1997b). *Readings on Color, Volume 1: The Philosophy of Color*. MIT Press, Cambridge, Massachusetts.
- Byrne, A. and Hilbert, D. R. (2003). Color realism and color science. *Behavioral and Brain Sciences*, **26**(1), 3–64.
- Byrne, A. and Hilbert, D. R. (2004). Hardin, Tye, and color physicalism. *The Journal of Philosophy*, **CI**(1), 37–43.
- Churchland, P. M. (2007). On the reality of objective colors. In J. Cohen and M. Matthen, editors, *Color Ontology and Color Science*. MIT Press, Cambridge, Massachusetts. Forthcoming.
- Cohen, J. (2003a). Color: A functionalist proposal. *Philosophical Studies*, **112**(3), 1–42.
- Cohen, J. (2003b). Perceptual variation, realism, and relativization, or: How I learned to stop worrying and love variations in color vision. *Behavioral and Brain Sciences*, **26**(1), 25–26.
- Cohen, J. (2004). Color properties and color ascriptions: A relationalist manifesto. *The Philosophical Review*, **113**(4), 451–506.
- Cohen, J. (2007). Color, variation, and the appeal to essences: Impasse and resolution. *Philosophical Studies*, **133**(3), 425–438.
- Hardin, C. L. (1988). *Color for Philosophers: Unweaving the Rainbow*. Hackett, Indianapolis.
- Hardin, C. L. (2003). A spectral reflectance doth not a color make. *The Journal of Philosophy*, **100**(4), 191–202.
- Hardin, C. L. (2004). A green thought in a green shade. *Harvard Review of Philosophy*, **XII**, 29–39.
- Hardin, C. L. (2007). Comments on Cohen’s ‘It’s Not Easy Being Green’. 2007 Society for Philosophy and Psychology, Toronto.
- Hilbert, D. R. (1987). *Color and Color Perception: A Study in Anthropocentric Realism*. CSLI, Stanford.
- Hubel, D. H. (1988). *Eye, Brain, and Vision*. Scientific American Library, New York.

- Hurvich, L. M. (1981). *Color Vision*. Sinauer Associates, Sunderland, Massachusetts.
- Hurvich, L. M. and Jameson, D. (1957). An opponent-process theory of color vision. *Psychological Review*, **64**, 384–403.
- Jakab, Z. and McLaughlin, B. (2003). Why not color physicalism without color absolutism? *Behavioral and Brain Sciences*, **26**(1), 34–35.
- Johnston, M. (1992). How to speak of the colors. *Philosophical Studies*, **68**, 221–263. Reprinted in Byrne and Hilbert (1997b), 137–176.
- Matthen, M. (2005). *Seeing, Doing, and Knowing: A Philosophical Theory of Sense Perception*. Oxford University Press, Oxford.
- McDowell, J. (1985). Values and secondary qualities. In T. Honderich, editor, *Morality and Objectivity: A Tribute to J. L. Mackie*, pages 110–129. Routledge and Kegan Paul, London.
- McGinn, C. (1983). *The Subjective View: Secondary Qualities and Indexical Thoughts*. Oxford University Press, Oxford.
- McLaughlin, B. (2003). The place of color in nature. In R. Mausfeld and D. Heyer, editors, *Colour Perception: Mind and the Physical World*, pages 475–502. Oxford University Press, New York.
- Mothersill, M. (1984). *Beauty Restored*. Oxford University Press, Oxford.
- Peacocke, C. (1984). Colour concepts and colour experiences. *Synthese*, **58**(3), 365–81. Reprinted in Rosenthal (1991), 408–16.
- Rosenthal, D. (1991). *The Nature of Mind*. Oxford University Press, New York.
- Sayre-McCord, G. (1988). The many moral realisms. In G. Sayre-McCord, editor, *Essays on Moral Realism*, pages 1–23. Cornell University Press, Ithaca.
- Smith, A. D. (2002). *The Problem of Perception*. Harvard University Press, Cambridge, Massachusetts.
- Sober, E. (1982). Realism and independence. *Noûs*, **16**, 369–385.
- Thompson, E. (1995). *Colour Vision: A Study in Cognitive Science and the Philosophy of Perception*. Routledge, New York.
- Thompson, E., Palacios, A., and Varela, F. (1992). Ways of coloring: Comparative color vision as a case study for cognitive science. *Behavioral and Brain Sciences*, **15**, 1–74.
- Travis, C. (2004). The silence of the senses. *Mind*, **113**(449), 57–94.
- Tye, M. (2000). *Consciousness, Color, and Content*. MIT Press, Cambridge, Massachusetts.
- Zangwill, N. (2001). *The Metaphysics of Beauty*. Cornell University Press, Ithaca.