

Redness, Reality, and Relationalism: Reply to Gert and Allen

Jonathan Cohen*

I will respond to some of their criticisms, but I want to say at the outset that in my view, when you publish something, you put it out there for all eyes to see from their own perspective. And so, I believe that all criticisms are fair. Not all are correct. But all are fair. We live in a Rashomon world (Kaplan, 2011, 504–505).

I owe thanks to Joshua Gert and Keith Allen for their thoughtful critical responses to my book, *The Red and The Real: An Essay on Color Ontology* (henceforth, *RR*). In this essay I respond to what I take to be the most important objections they raise.

1 Gert: Color Relationalism and the Multiple Aspects View

Although there is much on which Gert and I disagree, we are united in taking seriously as a constraint on theories of color ontology (as many have not) the significant variation in the way visual systems respond to one and the same color stimulus under different circumstances. Furthermore, we agree that much of this variation is blameless in the sense that (at least in many cases) no one variant should be theoretically distinguished from all the others as being uniquely veridical. That said, we disagree about the scope of application of such considerations about perceptual variation with respect to color, and about what conclusions should be drawn from them.

With respect to scope, I urge in *RR* that instances of perceptual variation arise in what appears to be a structurally identical form in intrapersonal cases (a single visual system responds differently to a single object in different perceptual conditions), interpersonal cases (the visual systems of different

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

people respond differently to a single object in the very same perceptual condition), and interspecies cases (the visual systems of organisms from different species respond differently to a single object in the very same perceptual condition). Given this apparent uniformity, I advocate a treatment of perceptual variation that extends smoothly to all such instances, and take it to be a benefit of the relationalist view about color defended in that work that it does extend — in what I hope is a theoretically satisfying way — to all these sorts of structurally similar variation. The key relationalist move, in a nutshell, is that we can accommodate the contemplated forms of variation by treating colors as relations to perceivers and perceptual circumstances. Thus, just as we can reconcile something's being your sister but not mine by treating *sisterhood* as a relation to individuals, we can reconcile something's looking red to you in your circumstance but looking not red to me in my circumstance by treating *redness* as a relation to perceivers and circumstances. Thus, for the color relationalist, the varying responses of visual systems to a single stimulus are states that represent distinct fine-grained relational properties such as *red for S_1 in C_1 , green for S_2 in C_2 , etc.*

In contrast, Gert is relatively unmoved by the apparent structural similarity of the different forms of variation, and appeals to different strategies to account for them. He proposes that his frankly hybrid approach to perceptual variation is, on balance, more plausible and less revisionary than relationalism, and therefore should be preferred to the latter on general grounds of rational conservatism even if, as he grants, there is nothing internally inconsistent or otherwise incoherent about the relationalist picture.

Below I'll consider Gert's objections against relationalism, and then assess the cluster of alternative strategies he offers to put in its place. I'll argue that relationalism is in better shape, and that his alternative approaches to perceptual variation are in worse shape, than Gert allows.

Before I begin, here are two remarks about objections I won't be making. First, in what follows I won't be plumping for my own account of variation over Gert's on grounds of theoretical uniformity. I agree with his contention that uniformity is a merely *prima facie* virtue that can be trumped by other considerations. While I think uniformity can be helpful in choosing between accounts that are otherwise very close competitors, there are in my opinion sufficiently many other differences between Gert's view and my own that there's no need to appeal to the merely *prima facie*, defeasible consideration of theoretical uniformity to choose between them. Accordingly, I'll put uniformity aside. Second, in order to keep the discussion manageable, I'll also mostly (except for note 21) ignore issues about interspecies variation; instead, I'll follow Gert in concentrating mainly on intrapersonal and interpersonal variation.

1.1 Relationalism and Intrapersonal Variation

1.1.1 An Objectionable Revisionism?

Gert's first reason for rejecting relationalism as an account of intrapersonal perceptual variation is his allegation that (particularly compared to the multiple aspects view he advocates; see below) color relationalism is objectionably revisionary — that it is “radical,” “surprising,” and “a little crazy.”¹ He writes,

I think it is safe to say that relationalism is not the preferred view of the philosophically or scientifically unsophisticated. Nor is it the preferred view of the philosophically sophisticated. After all, even in the philosophical literature, the claims that snow is white, blood red, and grass green are often taken as staple examples of the simplest, most basic and uncontroversial truths (2).

I believe this overestimates the revisionism of color relationalism.

It is indeed plausible that color relationalism, as developed in *RR*, goes beyond what common sense and untutored intuition tell us directly about color. This is, I take it, mostly what because common sense and untutored intuition tell us directly is much less than what it takes to spell out an adequate account of the metaphysics of color (or of much of anything). On the other hand, it is much less obvious whether and how relationalism says things that *conflict* with the deliverances of common sense.² Now, some have argued that relationalism conflicts with the broadly manifest image of color in some more indirect way — e.g., by making false predictions about the phenomenology of color experience or about how we ascribe colors to objects in natural language. I take these worries seriously, and devote part II of *RR* to showing that these ostensible conflicts are merely apparent — that the relationalist can explain, or explain away, the evidence that is claimed to tell against it. But Gert doesn't object to any of those proposed explanations on behalf of relationalism, so I doubt that this is the source of his dissatisfaction.

It is also true that philosophers not addressing themselves directly to issues about color use ordinary color ascriptions as examples of allegedly basic truths (which, according to my version of color relationalism, they are not). But it is not clear that this reflects any widely held deep anti-relationalist commitment as opposed to a provisional assumption inspired by the surface simplicity of (some) color predications — a *pro tanto* assumption made mainly in order to

¹I am gratified that he also calls the view “philosophically interesting” and “cool”. Unfortunately, in the context in which they occur, these are pretty clearly not intended as desirable attributes. Ah, well.

²To be sure, some philosophers have claimed (mostly without evidence) that naive intuition is directly at odds with the relationalist metaphysics of color. However, when investigators have actually solicited naive intuitions about the question of how colors are metaphysically constituted they have found that the latter are considerably more open to color relationalism than many have claimed (Cohen and Nichols, 2010). See also Allen (2012), who claims that “relationalism (or something approximating it) is a commonly held philosophical position amongst undergraduates” (11).

have some or other working example of a simple property.³ On the other hand, if there is a received view of the metaphysics of color among philosophers, that view is arguably some kind of dispositionalism — roughly, the view that colors are dispositions of objects to affect certain kinds of perceivers in certain kinds of conditions. For example, many find this view endorsed in the writings of the great modern philosophers (e.g., Galileo, Boyle, Newton, and Locke); more recent dispositionalists include McGinn (1983); Peacocke (1984); Johnston (1992). But since dispositionalism is a species of relationalism, there is a reasonable case to be made that the received philosophical view about color *is* relationalist.

I don't place much argumentative importance on these sociological observations, but do take them to suggest that relationalism is not quite so at odds with either common sense or philosophical tradition as Gert claims.

1.1.2 Uniformity and The Pink Shirt

Gert's main substantive objection to relationalism is that that view describes poorly our experience of a materially uniform pink shirt whose color appearance systematically changes across its spatial extent — perhaps it is brighter in a region closer to the room's light source, and darker in a fold where it is shadowed. Gert wants to say that this shirt is one uniform color over its entire spatial extent, and is worried that relationalism does not easily make room for this description. The worry is that, precisely because the relationalist takes the spatially varying color appearances to reveal distinct relational colors that the shirt genuinely exemplifies, she "must also say that talk of *the* color of the shirt, with the uniqueness the definite article implies, cannot be taken literally" (4). Gert recognizes that relationalism comes with two kinds of resources relevant to the description of the pink shirt, though he doesn't explain the relation between them, and instead treats them as independent lines of possible response to the case. In any event, he finds both of these lines of thought unsatisfactory, and so concludes that relationalism is inadequate to the case.

The first relationalist resource Gert considers is addressed to the intuition that the visually distinguishable regions of the shirt are visually represented as being in some sense uniform or stable. The relationalist account of this stability intuition rests in the observation that, though the distinct regions may occurrently exemplify different relational colors, they satisfy the following counterfactual: if, counterfactually, they were presented under the very same perceptual conditions, they would be a perceptual match (for certain kinds of visual systems, such as those of many ordinary trichromatic human beings). Hence, the many occurrently visually distinguishable regions of the materially uniform shirt all share the common — viz., stable across the shirt's spatial extent — property of being such that they would be perceptual matches were they presented under a fixed set of perceptual circumstances.

³I agree that if color properties are relational, this relationality does not show up explicitly in the surface form of (many) color predications; for discussion, see *RR*, chapter 4.

Gert accepts that there is this spatially stable property of the shirt's regions, and that it is visually represented. But he rejects this as an alternative account of the stability intuition on the grounds that it depends, ultimately, on something the relationalist rejects — the representation of a uniform color. Gert allows that at least some of the mechanisms contributing to color constancy might license representation of the counterfactual property to the regions without assigning a uniform categorical color. For example, he considers a strategy involving comparison of cone-excitation ratios of affected cone types across a luminance edge; the thought is that if the ratios are similar (although absolute excitation levels may differ), then this provides evidence that the counterfactual property holds between two regions, and does so without depending on assigning a common categorical color to the regions. But since, he says, this mechanism is at best one limited component of the mix of distinct mechanisms contributing to color constancy, it cannot plausibly be extended to all the cases where the stability intuition holds. And he thinks that there is only one explanation of the counterfactual property that *does* extend to all the cases where the stability intuition holds: the visual representation of the distinct regions as exemplifying a shared categorical color. But, of course, the relationalist cannot allow that. In short, then, the worry is that the relationalist's attempted explanation of the spatial stability of the pink shirt (in terms of the counterfactual property) must ultimately depend on the recognition that the shirt exemplifies a uniform color throughout its spatial extent, *pace* relationalism.⁴

This objection is unpersuasive for several reasons. First, while I sympathize strongly with Gert's claim that color constancy is an interaction effect (Cohen, 2012), Gert's objection turns on estimates about the range of cases to which different contributory mechanisms are applicable that, given the present state of understanding, we're not in a position to make with any confidence. Second, notwithstanding the last point, the range of cases for which there are relationalist-friendly explanations of the counterfactual property is considerably broader than Gert suggests, if only because (despite his framing of the issue exclusively in terms of cone excitation ratios) there are a variety of current proposals about the mechanisms subserving the computation of color constancy that work (or can be made to work) without prior assignments of uniform properties to surfaces.⁵ Given this point, one might reasonably

⁴Crucially, Gert is not objecting (as is Tye (2012), for example) that relationalism is unable to provide a story about the stability intuition. Rather, he is objecting that the story that (he allows) it *does* provide turns out to depend on mechanisms that cannot be understood other than in terms of the assignment of a uniform color.

⁵For example, there are models that begin from assumptions about the kinds of surfaces corresponding to median/extremal lightness values, and compute from these models of the incident illumination (Land and McCann, 1971; Buchsbaum, 1980; Gilchrist *et al.*, 1999). Similarly, there are views that attempt to recover a model of the illumination from information about mutual reflections in the scene (Funt *et al.*, 1991), the boundaries of regions known to be specular reflections (D'Zmura and Lennie, 1986; Lee, 1986), and shadows (D'Zmura, 1992). And there are proposals that appeal directly to assumptions about higher order scene statistics, (Golz and MacLeod, 2002; MacLeod, 2003; Brainard *et al.*, 2006). My point is (of course) not that any of these accounts is

be optimistic about the range of cases to which this sort of model applies. Third, and most significantly, it is simply not true that, in order to account for the stability intuition, relationalists must forswear appeal to representation of *any* uniform property of regions they claim are distinct in color. In order for an explanation of the stability intuition to abide by the relationalist's commitment that such regions differ in color, it only must be true that any such shared properties appealed to in the explanation *are not colors*. This matters because computational models of color constancy don't — indeed, can't — by themselves come with a (motivated) reason for identifying one sort of property as a color.⁶ For that you need a metaphysical account of color. Therefore, in the context of the current dispute about what the colors are, if it should turn out (for whatever reason) that the relationalist needs to say that two regions she distinguishes in color share a common property, that doesn't conflict with her relationalism. In light of these considerations, I don't see that Gert's objections show any substantial weakness in the relationalist's account of the stability intuition.

The case of the pink shirt suggests another intuition that makes for a *prima facie* challenge to relationalism; this is the intuition that what the visual system represents about the color of the shirt is not just stable/uniform, but also unique (hence, as Gert points out, we freely talk about *the* color of the shirt). Of course, relationalism entails that the shirt has a large (perhaps infinite) number of fine-grained relational colors, and so appears to conflict directly with the uniqueness intuition. However, the relationalist can resolve this apparent conflict by reference to the idea (discussed in *RR*, chapter 4) that the colors we ordinarily think and talk about — and the referents of the not-explicitly-relativized color predicates of natural languages — are relatively coarse-grained relational properties relativized to contextually salient perceivers and contextually salient viewing conditions. On this view, an utterance of '*a* is pink' in context *K* attributes to *a* the property *pink for the perceivers relevant in context K under the perceptual circumstances relevant in K* (*RR*, 100). In saying that Gert's imagined shirt exemplifies this coarse-grained relational property, the relationalist need not give up her claim that the shirt additionally exemplifies a large number of fine-grained relational properties (*pink for S₁ in C₁, red for S₂ in C₂, etc.*), Rather, she claims that there is a further relational color property — one that is the target of much of our ordinary thinking and speaking about color, that the shirt exemplifies as well. And I take it that the shirt plausibly does exemplify the latter property relative to the context in which Gert wanted to attribute pinkness to it — otherwise he would not have felt confident in saying that it is uniformly pink without further qualification.

Crucially, the relationalist will hold that the shirt's coarse-grained relational property is not just one among its many colors; rather, she will say that its

adequate to the range of cases by itself, but that the menu of options is wider than Gert's discussion suggests.

⁶This is so even conceding (as I am prepared to do for the sake of argument) a fairly robust sort of psychological realism about the computational models under discussion, without which Gert's objection doesn't get off the ground.

coarse-grained color occupies a special role in our thinking about color because this property is (as the others are not) relativized to the kinds of perceivers and perceptual circumstances relevant in the context of attribution. And this gives the relationalist materials to account for the uniqueness intuition. Her story will be that, even if there is not just one unique color that the shirt exemplifies, we can sensibly use the expression ‘the color of the shirt’ to pick out the one unique color relativized to contextually relevant perceivers and contextually relevant circumstances. (Analogy: although there is not just one unique remaining beer in the universe, we can sensibly use the expression ‘the remaining beer’ to pick out the unique beer remaining in the contextually relevant refrigerator.)

While it seems to me that this explanation accounts well for the uniqueness intuition, Gert disagrees, and objects that the proposed account overgeneralizes:

...even a patchwork shirt could be uniformly pink in this sense [viz., that of exemplifying the coarse-grained color *pink for the contextually relevant perceivers under the contextually relevant perceptual circumstances*], as long as each of the slightly differently colored pink patches would appear to be some shade of pink to normal viewers. But the shirt of our example is not like this (5).

But this objection is puzzling. Gert is right that the shirt of his initial example and the patchwork shirt differ: the former but not the latter is chromatically stable or uniform (in the sense of the stability intuition discussed above). However as discussed, the relationalist can agree that the initial pink shirt is and the patchwork shirt is not stable, and indeed has an account of this difference (given in terms of counterfactuals relating distinct regions of each shirt). Admittedly, the case might make trouble for the relationalist if we had reason for believing that whatever relationalists say to account for the uniqueness intuition must also serve as an account of the stability intuition. But we don’t, so it doesn’t.

1.2 Gert on Intrapersonal Variation

As noted, Gert criticizes the relationalist account of intrapersonal variation as part of a comparative case for what he claims is a more conservative “multiple aspects” view, on which (non-relational) colors can present (veridically) a variety of aspects, appearances, or ways/modes of appearance under different conditions:

On this view colors can present different aspects in different viewing conditions, just as shapes can. ...it should be apparent to readers that it can be obvious that a surface is uniform in color even though the appearance of the surface changes quite considerably from region to region — as in the example of the pink shirt (4).

By holding that distinct aspects of this kind can non-competitively present colors, Gert hopes that his view can, like relationalism, avoid the difficulty of having to single out one perceptual variant as being uniquely veridical. But he hopes the view can secure this benefit without committing to what he regards as the revisionary excesses of relationalism.

When considered so schematically, Gert's guiding idea that colors present multiple appearances is attractive. However, it is far less clear just how this idea should be understood more precisely, and whether it can be extended into a coherent view that is both plausible and an alternative to color relationalism.⁷

Putting that point aside, one pressing worry about the proposal is that it seems to mislocate the subject matter of ordinary color discourse.⁸ Thus, an observer who claims that an a 510nm spectral light is unique green appears to be, and is ordinarily described as, reporting on the light's *color*.⁹ But on the multiple aspects view, the state/report is not about the light's color, but about the way its color is presented.¹⁰

⁷In particular, it is worrying that the view helps itself to a technical notion of "aspects" or "appearances" of color properties that cries out for further elucidation. Initially, one might hope to dispel the mystery surrounding the relevant "aspects" by identifying them with some sort of perceptual modes of presentation of properties — roughly, perceptual analogues of Fregean senses for perceived objects; unfortunately, however, there is deep controversy surrounding the comparatively clear case of the interpretation of Fregean senses in language, and even less consensus about whether and how the apparatus might be extended to perception. Gert acknowledges this, but tells us that he is content to fall back on a "less technical interpretation of 'mode of presentation', on which it means 'way of appearing'" (8).

Fair enough. But it is not obvious that the thin notion of ways of appearing found in ordinary discourse is ultimately coherent or compatible with other well-justified commitments (this is something we would expect to see worked out in a theoretical elaboration, not in common usage itself). Moreover, I doubt that the common sense notion is sufficiently committal to guarantee that distinct ways of appearing can avoid conflict, as Gert needs them to: after all, sometimes we do say that the way something appears to me now conflicts with the way it appears to me later. As things now stand, then, the alternative account Gert offers turns on a piece of apparatus that goes beyond what common sense can underwrite, and for which he demurs from offering further explanation. I suggest that this renders the proposed alternative less than satisfying.

⁸In assigning contents to both perceptual experiences and reports about them, here and below, I am assuming that perceptual experiences determine correctness conditions. (For defense of this assumption, see Byrne (2010); Pautz (2010); Siegel (2010); for criticism, see Travis (2004)). I am not assuming that these correctness conditions can only be met by the obtaining of an "isomorphism" between the experience and the distal property that is its content, as Gert speculates at one point (17).

⁹Unique hues are typically described as those that are "perceptually unmixed". Normal visual perceivers will make relatively intrapersonally stable choices of stimuli (e.g., spectral lights or Munsell chips) that look unique green (i.e., the sample looks green without looking at all bluish or at all yellowish), unique red (looks red without looking at all bluish or yellowish), unique blue (looks blue without looking at all reddish or at all greenish), and unique yellow (looks yellow without looking at all reddish or at all greenish). The locus classicus for discussion of unique hues and their import to philosophical accounts of color is Hardin (1988).

¹⁰In making this objection (in anticipation of a version of Gert's view) in *RR*, I wrote that such a report "would ordinarily be characterized as a report of the *color* experienced by the subject — not the way in which the color is presented" (91). Gert finds that the phrase 'the color experienced by the subject,' as it occurs in the formulation of the objection, is ambiguous between a reading that picks out a way of appearing and one that picks out an "objective color" (9). For what it's worth, I don't believe the phrase in question is ambiguous — I don't hear any ambiguity, and I believe the phrase fails standard tests for ambiguity. However, if it is ambiguous, I certainly didn't intend

Gert responds to this objection by urging that, in fact, sometimes we *do* seem to make reports about how things look in respect of color that are best understood as having contents about the way colors appear rather than colors, hence that his view is not guilty of any serious mislocation. Here is his example:

If I am looking at an object that appears to me to be uniformly colored green, but that happens to be situated near a blue wall, I might utter the following three claims.

(U3) The object is manifestly a uniform shade of green.

(U4) There [demonstration of a part of the object far from the wall] it looks unique green.

(U5) There [demonstration of a part of the object near to the wall] it looks bluish green.

Given my expressed view in (U3) that the object is uniform in color, the natural characterization of (U4) and (U5) cannot be as reports of the objective color the object appears to have (9).

But this example doesn't show what Gert takes it to. As far as I can tell, (U4)/(U5) have two readings, corresponding to different disambiguations of their occurrences of 'it', but neither of which sustains Gert's contention.¹¹ On the first, 'it' is anaphoric on 'the object' in (U3), so (U4)/(U5) assert that, in one of its demonstrated spatial parts, the *object* is visually represented as bearing the property *unique green/bluish green*. Alas, this reading won't serve Gert's purposes, since it construes (U4)/(U5) as reporting facts about the way the object looks, not the way any color looks. The second reading, on which 'it' in (U4)/(U5) is anaphoric on (U3)'s occurrence of 'a uniform shade of green', is *prima facie* more promising for Gert, since on this reading it may seem that the attributions have features of a color as their contents. On *secunda facie*, however, the imagined support this reading gives for Gert's view falls away. Again, the flatfooted (and now standard) understanding of such looks locutions — viz., those of the form '*P* looks *F*', with predicates for colors or other sensible properties as their grammatical subjects — treats them as amounting to something like '*P* is visually represented as being *F*', or '*P* is

the objection to be construed as on Gert's first reading: the point of the objection was to urge that the ordinary report is naturally interpreted as *not* concerning the way of appearing. Therefore, in reformulating the mislocation objection in the main text I have re-expressed it so as to avoid any such ambiguity, and in the discussion that follows I consider only what Gert says about the second reading.

¹¹Gert could avoid the pronoun-binding ambiguity by use of the following alternative formulations:

(U4') There [demonstration δ] the color looks unique green.

(U5') There [demonstration δ'] the color looks bluish green.

These formulations can only be read in a way analogous to the second reading of (U4)–(U5) considered below; however, I'll argue that the example doesn't go through on this reading, so the contemplated change doesn't help.

visually represented to me as if it were F' .¹² But if so, then on the reading under consideration, (U4)/(U5) are reports about which color the object is visually represented as being — which is just to say that, on this reading, too, they *are* reports about the color the object has, contrary to what Gert claims.

Now, in the passage quoted above, Gert claims that his conclusion about the content of (U4) and (U5) follows from their interaction with (U3); and I have made only passing use of (U3) so far. Do things change if we suppose that the subject additionally visually represents (U3)?¹³ I don't see that they do. For this supposition could potentially affect the contents of (U4)/(U5) only if the subject's visually representing the object's color as uniform somehow prevented her from having the variety of visual representations of its color reported by (U4)/(U5). But, it's hard to see why this should be the case. For one thing, famously, a subject's visual representations don't all have to be consistent, so it's not clear why there should be the contemplated kinds of exclusion relations between visual representations. For another, as discussed in §1.1.2 above, there are several ways of understanding the content of (U3) that render it consistent with (U4) and (U5) (where the latter are taken as reports of varying visual representations of the color of an object).

In short, then, appeal to (U3) gives us no reason for wavering from the understanding of the contents of (U4)/(U5) on which the latter are reports about the colors of objects. Since the multiple aspects view entails that these — and all similar reports — are not reports about the colors of objects, the charge that that view suffers from a mislocation problem stands.

1.3 Gert on Interpersonal Variation

Although, as noted, Gert proposes different approaches to intrapersonal and interpersonal instances of perceptual variation, there are important structural similarities between his different approaches. In both sorts of cases, his view is that there is no blameless variation with respect to the representation of *color*, but only variation with respect to something that is finer-grained than color. On the intrapersonal side, we have seen that the finer-grained entities in question are *aspects* of colors. On the interpersonal side, he holds that the finer-grained entities in question are relatively determinate *precisifications* of colors.¹⁴ In both sorts of cases, his view is that experiences of color (for a materially

¹²This is the comparative, or epistemic use of 'looks' discussed by (Chisholm, 1957; Jackson, 1977).

¹³The quoted passage doesn't specify whether (U3)'s content (that the object is manifestly uniform) is supposed to be visually represented, or merely believed/asserted. But if it is not visually represented by the subject, then there's no reason to take it as constraining the contents of the visual representations she reports by (U4) and (U5) — this is why, for example, visual illusions persist even after we recognize that they mislead.

¹⁴Gert's way of putting this point involves holding that colors are "vague properties" (12), which suggests that he's assuming (controversially) that the relevant vagueness qualifies the properties as opposed to the predicates. But I believe he can capture much of what he wants to say without taking any stand on controversial issues about vagueness by instead formulating the point in terms of the relations of relative determinacy between properties. His view, then, would be that colors are less determinate than their precisifications, and that extensive blameless interpersonal perceptual

uniform single stimulus) must be unvarying or uniform to be veridical, so there can't be blameless perceptual variation with respect to colors; but he allows that there can be blameless perceptual variation with respect to something finer-grained (intrapersonally: aspects/ways of appearing; interpersonally: precisifications/determinates) such that there's a one-many mapping from the former to the latter. And his strategy, in both instances, is to insulate colors from considerations about blameless perceptual variation by thinking of all the variation as happening exclusively at the non-color level.

To see how this approach works concretely, consider the form of interpersonal variation that has figured most centrally in recent philosophical discussion of color — that involving the finding that settings for the four chromatic unique hues vary significantly among visual perceivers who pass standard clinical tests for normal color vision. Gert accepts that there is this sort of extensive interpersonal variation, and, moreover (unlike many), that it is blameless in the sense that there doesn't seem to be anything that could possibly make it the case that one normal perceiver is correct — to the exclusion of others who disagree with her — about whether a 510nm spectral light, as it might be, is really an instance of *unique green*. But since he construes the variation in question only as variation about the “way a colored object can appear” (10), he denies that these claims reveal any blameless variation with respect to *color*.¹⁵

Gert admits that it is a counterintuitive consequence of his approach to interpersonal variation that it requires saying that *unique green* is not a color, but he regards this as a price worth paying. Indeed, he claims, the existence of extensive blameless perceptual variation with respect to *unique green* (and the like) is itself sufficient reason for denying that the latter is an object color:

a natural response to the question ‘Where, given this variability, is unique green *really* to be found on the spectrum?’ would be ‘Nowhere’. And if it is really to be found nowhere, it is not an object color (10).

if there is sufficient interpersonal variability in color experience when those experiences are described with a high degree of precision (‘unique green, ‘perfectly balanced orange’), then there is no fact of the matter regarding the stable color of that object that has that degree of precision (10).

Unfortunately, cases of extensive interpersonal perceptual variation extend more broadly than Gert allows; consequently, consistent application of his view results in denying the status of colors to not only fine-grained determinate

variation about the chromatic is confined to the relatively determinate level, so doesn't apply to colors properly speaking.

¹⁵Kuehni (2004) provides a useful overview of the extent of interpersonal variation for unique hues. For alternative philosophical responses to this sort of variation that turn on denying that the variation is blameless, see Tye (2006a,b, 2007); Byrne and Hilbert (2007); but see Cohen *et al.* (2006a,b).

properties like *unique green*, but also to less determinate, paradigmatic color properties — those lexicalized in English color terms — such as *green*, *blue*, *orange* and *red*.¹⁶ One impressive demonstration of such variation with respect to intermediate/lexical colors comes from Malkoc *et al.* (2005). These authors describe a color naming experiment in which normal subjects were shown stimuli spanning the color circle in roughly equal steps, and were asked to label the stimuli using one of eight color terms: four corresponding to the unique hues ('red', 'green', 'yellow', and 'blue') and four corresponding to the binary hues ('orange', 'purple', 'blue-green', and 'yellow-green'). It turns out that there was extensive disagreement in this naming task. For example, they report that seven hue samples were labeled 'blue' by at least one subject, but there was 80% consensus in the application of that label to only two (and that all seven were assigned color labels other than 'blue' at least once). Similarly, six samples were labeled 'yellow' by at least one subject, but no stimulus commanded more than 60% consensus in meriting the label 'yellow'. (Examination of figure 4 of Malkoc *et al.* (2005) shows that this degree of non-overlap was characteristic of all of the eight unique and binary color categories). On the evidence, then, there appears to be just the same sort of extensive and apparently blameless interpersonal variation with respect to the representation of blue, yellow, and the rest of the intermediate-level color properties that there is for the unique hues. Again, given his other commitments, this suggests that Gert will be forced to say that red, green, purple, yellow green, blue, yellow, blue green, and orange are not colors. That seems bad.

Gert is aware of the threat to his view by such results, and offers three objections in an attempt to explain away the recalcitrant evidence. First, he follows (Byrne and Hilbert, 2007, note 5) in worrying that the color-naming protocol and large within-subject variability reported by Malkoc *et al.* (2005) complicates the interpretation of their results. But this concern is dubious:

¹⁶A similar worry can be raised with respect to the superdeterminable properties *bluish*, *yellowish*, *greenish*, and *reddish* in terms of which the unique hues are defined. For if subject S_1 represents a 510nm spectral light as unique green and S_2 represents a 510nm spectral light as bluish green, *ipso facto* they are manifesting interpersonal variation with respect to *bluish*. And if the variation with respect to the unique hues is blameless, then so too must be the concomitant variation with respect to the superdeterminables. It might seem then, that because Gert is willing to infer from interpersonal variation data that the unique hues are not colors, he is likewise committed to saying that *bluish*, *yellowish*, *greenish*, and *reddish* are not colors.

But Gert can avoid that consequence. For interpersonal variation in the unique hues only commits Gert to saying that there are *some* cases of blameless interpersonal variation with respect to the superdeterminables. And if it turns out that, despite these cases, there is, overall, substantial interpersonal consensus about the application of these superdeterminables (i.e., that, overall, interpersonal variation with respect to these superdeterminables is not blameless), then he won't be committed to denying their status as *bona fide* colors/properties. The relevant distinction between *unique green* and *greenish*, then, is that, though there are cases of blameless interpersonal variation with respect to both, there seems to be nothing but blameless interpersonal variation (i.e., no consensus cases) for *unique green*, while Gert thinks there are sufficiently many consensus cases for *greenish*.

I suppose one could worry about how many cases of consensus, or what degree of consensus within a community, is necessary for genuine color properties (Gert doesn't say). In any case, I won't press these worries.

while I'm not sure just what about color-naming is worrying Gert (or Byrne and Hilbert), Malkoc *et al.* (2005) address the worry about large within-subject variation by separately analyzing settings for the most consistent subjects, and report that for these the correlations in repeated settings reached 0.8 (indicating between-subject variance roughly twice that of within-subject variance).¹⁷ Second, Gert objects that, in order to maintain constant luminance and saturation levels across stimuli, Malkoc *et al.* (2005) had to present samples of some stimuli that were not typical/good focal exemplars of their categories. Again, this concern is unpersuasive. It is true that some of the stimuli used in this work may not have represented ideal focal examples of their categories (in particular, for yellow because their foreground stimulus was isoluminant with the background, and for red because the saturation was somewhat low). However, the other six of the eight intermediate level properties tested are relatively independent of saturation and lightness, so there is no reason to suspect that their samples of the other colors were atypical. Since Malkoc *et al.* (2005) found significant interpersonal variation for all eight of the intermediate level colors tested, the non-focality of the samples corresponding to 'red' and 'yellow' can't explain the observed variation.¹⁸

Be that as it may, Gert's most important reason for wanting to dismiss the empirical evidence of significant interpersonal variation about the intermediate/lexical colors is not methodological, but transcendental. He holds that there *couldn't* be the sort of variation Malkoc *et al.* (2005) and others claim to have found, since extensive interpersonal variation of that sort would undermine the possibility of ostensive learning of color terms:

There has to be sufficient consensus on a sufficient class of objects, as to their broad colors, to enable children to learn terms for these colors by ostension and correction, since children obviously do learn these terms. Any child who does not see the fire engine in her "Babys First Color Book" as red, or the duck as yellow, or the sky as blue, is suffering from some visual defect (12).

Once again, this consideration is unconvincing. For one thing, though it is obvious that children do acquire color terms, it is not at all obvious whether and to what extent ostension and correction play a significant role in the process; as such, it is deeply unclear whether the variation in question conflicts

¹⁷Because Malkoc *et al.* (2005) don't report within-subject variation for the specific result at issue (that represented in figure 4), it is not clear what prompts the concern here. However, for other tasks on which they report, they address this issue indirectly by reporting correlations between repeated settings of the same hue: if all the variability were within-subjects, this correlation would take the value 0, while if all variation were between-subjects, the correlation would be 1. They report (2158) that such correlations were in all cases significant, which is to say that there was statistically significant between-subject variation.

¹⁸I would add that, though I follow Gert above in confining attention to the striking evidence of interpersonal variation in perception of lexical colors provided by Malkoc *et al.* (2005), this ignores the independent (cross-cultural) empirical support for the same conclusion provided by, e.g., Webster and Kay (2006); Lindsey and Brown (2009), to which Gert's present methodological worries are not applicable. (I am indebted to Mike Webster for discussion of the issues in this paragraph.)

with facts about learnability.¹⁹ For another, the kind of broad agreement in color term application that would presumably be required by uniform labeling of the objects in “Baby’s First Color Book” is not inconsistent with interpersonal variation in the application of intermediate level color terms. We know this must be so because even anomalous trichromats, dichromats, and others who fail standard tests for normal color vision nonetheless largely converge on the same application of color terms as those who do pass those tests; this is why, for example, their visual anomalies typically go undetected for years. *Inter alia*, this means that, despite varying extremely significantly from normal perceivers in their perception of intermediate/lexical colors, these perceivers perform like the rest of us with “Baby’s First Color Book”.²⁰ The upshot, then, is that, while the shared enjoyment of “Baby’s First Color Book” may require some amount of communicative convergence in the application of lexical color terms, this gives us no reason to doubt that (just as the data indicate) there is significant interpersonal perceptual variation with respect to the intermediate colors blue, yellow, green, orange, and the rest.

In summary, then, there appears to be significant and extensive interpersonal perceptual variation with respect to both fine-grained chromatic properties (such as *unique green* and *unique blue*) and intermediate/lexical chromatic properties (such as *red* and *orange*). In the face of these results, Gert’s strategy of describing variation as happening exclusively at the non-color level will require that all of those chromatic properties are not colors. I suggest that a theory of color that excludes from its domain not only *unique green*, *unique blue*, *unique red*, and *unique yellow*, but also *green*, *blue*, *red*, *yellow*, *blue green*, *yellow green*, *orange*, and *purple*, is not a theory we should want.²¹

¹⁹It is perhaps worth reminding ourselves that, over a long period, many psychologists and linguists assumed that all language learning worked through ostension and correction. Chomsky’s suggestion that it might not is a crucial element in his rejection of behaviorism, and therefore in the development of modern cognitive science.

²⁰Blatant conjecture: I suspect the observed convergence in labeling despite differences in perception can be explained (given the assumption of perceptual variation) in terms of the routine operation of conversational accommodation. (Analogy: although the NBA coach enters the conversation with his own range for the application of the gradeable adjective ‘tall’, he easily accommodates the high school coach’s application of the predicate to a high school player whose height is six feet.)

Whether this (entirely conjectural) explanation of the convergence is correct or not, the point in the main text stands: convergence in the communicative use of color terms is compatible with large interpersonal differences in the perceptual representation of (intermediate/lexical level) color categories.

²¹ Perhaps this is the place for a brief remark about Gert’s deflationary treatment of interspecies variation. Gert holds that if S_1 and S_2 are members of different species, the perceptual states of S_1 and S_2 can’t disagree (or agree) about whether o exemplifies P for any color P , since, roughly, no single color is in the representational repertoire of perceptual systems from distinct species. And he supports this view by claiming (14–15) that members of different species exhibit different psychological color spaces.

While this proposal merits more discussion than I can give it here, I’ll just note that, in my view, Gert is here both overestimating interspecies differences and underestimating intraspecies differences. Our discussion of interpersonal differences has already brought to light the wide range of at least certain kinds of (blameless) intraspecies variation; and while Gert may wish to insist that all such variation is not blameless but evidence of a defect (15), this reaction is both

2 Allen: Relationalism and Subject-Involving Contents

Keith Allen is interested in, and critical of, a specific commitment of color relationalism — viz., the claim that subjects are (always) constituents of the contents of perceptual, cognitive and linguistic representations of color.²² He offers different arguments against subject-involving color contents in perceptual representations on the one hand and cognitive and linguistic representation on the other. I'll reply to these considerations in §2.1 and §2.2, respectively.

2.1 Subject-Involving Perceptual Representations?

2.1.1 Motivation

Allen begins his discussion of the representation of color in perceptual experience by observing, plausibly, that subject-involving contents are more complicated than subject-independent contents. For this reason, he suggests that we need motivations for embracing the more complicated, subject-involving option.

I agree wholeheartedly with Allen that subject-involving perceptual representations of color require motivation, and devote a large chunk of *RR* to the task of offering such a motivation by reference to facts about perceptual variation. Here is a short version.²³

Consider an ordinary instance of perceptual variation — a case in which two different visual systems have psychophysically distinguishable experiential reactions to one and the same color stimulus, and let it be that (as per standard but not universally shared assumptions; see note 8) those reactions represent the color of the stimulus, or at the very least determine correctness conditions about the stimulus color. Then we can ask which (if any) of those representations is veridical. The following four answers exhaust logical space: neither, the first but not the second, the second but not the first, or both. But the neither-right answer is, plausibly, unduly skeptical: it leads pretty quickly to the view that we all suffer from systematic and widespread perceptual error as a matter of course. Moreover, the first-but-not-second and

stipulative and (on its face) inconsistent with the empirical evidence concerning normal human color perceivers discussed above. On the evidence, there appear to be significant differences between the color spaces of normal trichromatic human observers, greater differences between those of normal trichromatic human observers and anomalous trichromatic human observers, greater differences still between normal trichromatic human observers and dichromatic human observers, and so on. But it is unclear that there's the qualitative break between the intraspecies differences, on the one hand, and interspecies differences, on the other, that Gert seems to be assuming.

²²Allen argues against the commitment to subject-involving contents that he finds in both the relationalism of *RR* and also the “centering feature” view of Egan (2012); unfortunately, I must omit discussion of the latter here for reasons of space.

²³I present versions of this argument in several places; of course, my fullest elaboration and defense of the argument occurs in *RR* itself.

second-but-not-first answers are hard to motivate: since perceptual variation is compatible with making the representations symmetric in just about any way that could conceivably matter, it is hard to see what (beside ad hoc stipulation) would make it the case that one of them is veridical to the exclusion of the other. That leaves the final option — both representations are veridical — as the best reaction to cases of perceptual variation. But this ecumenical conclusion, in turn, motivates taking perceptual representation to be subject-involving, in so far as the latter view explains how it could be that both representations could be compatibly veridical. For the subject-involving color contents (Color_1) and (Color_2) can be compatibly veridical in exactly the way that the subject-involving contents about family members (Sister_1) and (Sister_2) can be compatibly veridical.

(Sister₁): Jane is the sister of S_1 .

(Sister₂): Jane is not the sister of S_2 .

(Color₁): Spectral light a is unique green for S_1 .

(Color₂): Spectral light a is not unique green for S_2 .

There is, of course, much more to be said about this argument than I can say here. Still, I hope that the abbreviated version presented here shows that the relationalist has something to offer in response to Allen's (reasonable) demand for a motivation for subject-involving contents.

Questions of motivation aside, Allen presents two further objections against subject-involving contents of perceptual experience. First, he worries that the subject-involving aspect of such contents is redundant. Second, he objects that the view overgeneralizes in such a way as to preclude perceptual error.

2.1.2 Redundancy

The worry about redundancy is that subject-involving content serves no purpose for visual systems: a subject perceives the world with the same visual system on every occasion, so specifying the subject in visual contents conveys no information, and can serve no adaptive purpose for the organism.

My reaction to this objection is to accept it. As noted above, the most important motivation for accepting subject-involving perceptual contents is that such contents serve a purpose for (not the organism, but) the theorist — viz., that of allowing her to describe perceptual experience in a way that avoids undue skepticism. The content of perceptual experience is, of course, a theoretical construct; as such, it should be characterized in whatever way best serves the explanatory needs it is enlisted to serve. If I am right that subject-involving contents serve the theoretical ends of explaining perceptual variation without committing to systematic skepticism better than subject-independent contents (and assuming that is a theoretical end worth having), then that should be reason enough for accepting subject-involving contents.

2.1.3 Lost Error

Allen's second objection is that making the contents of perceptual experience subject-involving threatens to foreclose the possibility of perceptual error. The worry is that the relationalist's ecumenicism overgeneralizes — that her attempt to avoid asymmetric errors in the perceptual representation of color amounts to holding that every (or nearly every) perceptual representation of color is veridical. But this is just to say that (nearly) none of those perceptual representations is erroneous.

Now, as Allen is well aware, representationalists have a strategy for answering the charge of lost errors. This strategy involves the idea that we have two separate levels of representation for color — a fine-grained, highly determinate perceptual representation, and a coarser-grained, less determinate representation for thought and language.²⁴ On this proposal, a perceptual experience might represent light *a* as exemplifying *red for S in C*, where '*S*' and '*C*' are relatively detailed specifications of my visual system and my perceptual circumstance. Nonetheless, an utterance/thought about the color of *a* occurring in context *K* would attribute to *a* a different and less determinate — as it might be, *red for the perceivers relevant in context K under the perceptual circumstances relevant in K*. This appeal to coarse-grained color representation helps in answering the charge of lost errors because it allows that there might be representational errors about color even if, per the relationalist's ecumenicism about perceptual representation, (nearly) all perceptual representations of color are veridical. For example, suppose *S*'s perceptual system veridically represents *a* as exemplifying *red for S in C*; still, if *S* wrongly believes that *S* and *C* are *K*-relevant perceivers/circumstances, she may infer/compute from her perceptual representation to the erroneous conclusion that *a* exemplifies *red for the perceivers relevant in context K under the perceptual circumstances relevant in K*. If she represents that conclusion, she will be erroneously representing *a*'s color.²⁵

²⁴This is just the proposal discussed as a way of responding to one of Gert's objections involving the uniform pink shirt in §1.1.2 above.

²⁵As Allen is also aware, this strategy is not the only means by which relationalists can allow for perceptual errors of color representation. Here are three other relationalist-friendly strategies that come to mind.

A first additional strategy involves hallucination. Thus, suppose, in a hallucinatory experience, I perceptually represent my hat as exemplifying *red (for myself in the circumstances I am in)*, although my hat is nowhere in the vicinity. Then the relationalist is within her bounds to call the hallucinatory experience an error involving the perceptual representation of color.

A second additional strategy involves the representation of the sorts of counterfactuals relating different colored objects/patches discussed in §1.1.2. To see what this amounts to, consider Adelson's well-known checkershadow illusion (http://web.mit.edu/persci/people/adelson/checkershadow_illusion.html), in which, given the illumination cues and other details about the configuration, viewers typically find it surprising to learn that two very different looking patches labeled '*A*' and '*B*' are intrinsically qualitatively identical. The relationalist will describe this case as involving an illusion about the color that the chips *would* be manifesting were they presented under the same viewing condition. She will say that (in addition to representing *A* and *B* as occurrently manifesting different colors) subjects represent *A* and *B* as being such as to mismatch perceptually were they

Unfortunately, Allen is dissatisfied with this answer. He worries that, though it may save the possibility of error for the relationalist, it doesn't allow for errors of *perceptual* representation, since it locates the error in the (post-perceptual) coarse-grained representations figuring in the cognitive and linguistic representation of color. He presses this worry by considering a grapheme→color synesthete in whom exposure to the numeral '6', printed in black ink, reliably causes (automatic/involuntary) representations of both green and black. Assuming that these representations are both perceptual, relationalists seem committed to saying that both are veridical, hence that the numeral *is* both black to this subject under this circumstance and green to this subject under this circumstance. "But," he writes, "this doesn't seem especially plausible" (5).

Alas, I find this objection less than convincing: I see three (independent) lines of response to it that are available to relationalists. First, I confess to not seeing what is problematic in holding that the numeral genuinely exemplifies both relational colors. (Surely the worry can't be that this would make those properties collapse; after all, the properties will come apart in their application to other objects even if they coincide in applying to inscriptions of '6'.) If so, then the case under consideration is not a case of misrepresentation of color, so not a lost error that relationalists are unable to describe. Second, the objection depends on presupposing that synesthetic color representations are perceptual; but this is deeply controversial, at least partly because there is no generally accepted method for distinguishing between perceptual and cognitive representations. But if the presupposition is false, then the case is ripe for treatment in terms of the relationalist strategy outlined above. Third, even those (such as myself; see Cohen (2013)) who regard synesthetic and normal perception as importantly similar will nonetheless allow that there remain respects in which synesthetic representations are *abnormal*, which suggests that their etiology is in some respect deviant (whence the standard

(as they are in fact not) presented under a shared circumstance — in effect, their visual systems are predicting that the chips would look different (would not match) were they presented in identical circumstances. But this last representation is erroneous: the regions would indeed be a perceptual match were they presented in identical circumstances. Thus, the relationalist has the means to say that, in this and other cases we normally describe as color illusions, the visual system is indeed misrepresenting a color property.

A third additional strategy involves deviant causal chains. Thus, to return to a fanciful example from Cohen (2007), if a telekinetic tomato stimulates my visual cortex without affecting my retina, and subsequently attending to that tomato causes in me the reaction I normally get when looking at ordinary, non-telekinetic, ripe tomatoes, then my perceptual representation of its color is obtained through a deviant causal chain. If this is the right description of the situation, then it seems plausible on grounds independent of relationalism to say that the perceptual representation in question is erroneous by virtue of having a deviant etiology. (Allen objects (5) that the admittedly unusual etiology in question still runs through the visual system, so is perhaps not deviant after all. In fairness, it is difficult to judge the particular case in the absence of a principled criterion of deviance; and I don't have one. But the viability of the strategy on offer is independent of verdicts about the particular example. The point is that a relationalist can — for principled reasons independent of her relationalism — describe as erroneous any perceptual representations of color whose etiologies count as deviant. As such, the existence of deviant perceptual causation offers the relationalist yet another way of describing some perceptual representations of color as erroneous.)

description of the condition as a psychological/neural pathology). But if so, then this very deviance is a reason for describing such representations as erroneous, even assuming the general relationalist view that *non-deviant* perceptual representations of color are almost always veridical (cf. note 25). Given the availability of these responses, I do not see that consideration of synesthetic color representations reveals any serious shortcoming in the relationalist's account of color misrepresentation.

2.2 Subject-Involving Cognitive/Linguistic Representations?

RR claims not only that perceptual representations of colors are subject-involving, but that the same is true about cognitive and linguistic representations of color. Once again, Allen disagrees: he holds that cognitive and linguistic color contents are subject-independent.

Now, it must be granted that ordinary color ascriptions in natural language — which presumably comprise a crucial source of evidence about the contents of linguistic and cognitive representations of color — are typically not overtly subject-involving. We ordinarily utter (1a) rather than (1b):

(1a) That ripe lemon is yellow.

(1b) That ripe lemon is yellow for *S* (in *C*).

But this does not establish that linguistic/cognitive color contents are not subject-involving, for it could be that subject-involving elements are added (possibly in a way that is sensitive to features of the context of utterance) to these contents despite not being present in the overt linguistic form of utterances expressing them. Something like this story is plausible for many different natural language expressions; for example, it is standard to hold that gradeable adjectives (e.g., 'large,' 'expensive,' 'tasty,' 'close') have as part of their content a covert/unpronounced, contextually determined, standard of comparison (large for an elephant, expensive for a taco, etc.).²⁶ In RR I propose that something similar is true about linguistic/cognitive color contents — viz., that the latter have a (contextually supplied) subject-involving element despite the absence of overt elements of this sort in the surface form of linguistic color ascriptions. And I offer a contextualist semantics for natural language color terms that is intended to sustain this suggestion. Specifically, I propose that the predicate 'is yellow', as uttered in context *K*, expresses the property *yellow for the perceivers relevant in context K under the perceptual circumstances relevant in K* (100; *mutatis mutandis* for other color predicates).

Allen argues that this proposal misdescribes the semantic behavior of ordinary color terms. For, he claims, such terms fail three standard diagnostics for context-sensitivity (compleatability, context shifting, and collection; see §2.2.2). He allows that one might nonetheless defend a revisionary contextualist semantics as a proposal for linguistic reform; but he takes the evidence to

²⁶See Klein (1991); Kennedy (2007).

show that the contextualist treatment is descriptively inadequate. And since he thinks contextualism remains the best hope for defending subject-involving cognitive and linguistic color contents, he concludes that such contents are, on the evidence, best construed as subject-independent.

2.2.1 Contextualism as Revision

I will argue (§2.2.2) that the evidence against the contextualist proposal for color predicates is indecisive. However, before I come to that, I want to concede that Allen's main conclusion may be true: I accept that contextualism may be descriptively inadequate as an account of ordinary color language. (My concession is not that the evidence demonstrates the descriptive inadequacy of the view decisively; it is that, for all the evidence shows, the view *may* be descriptively inaccurate.) Instead, I have come to think that the contextualist semantics in *RR* is best construed as a revisionary proposal.²⁷

To see what I have in mind, consider the analogous question about how we should think about the semantics of motion predicates after the discovery that, on the best, broadly empirical account, motion properties are constituted in terms of relations to reference frames. Of course, our motion predicates (e.g., 'is moving at 60mph') predate this broadly empirical discovery. As such, it is extremely plausible that these predicates do not come equipped with unfilled covert parameters for reference frames as part of their lexical information; and they are certainly not overtly relativized to reference frames. But if so, then it is hard to see how such predicates could have as semantic values the kinds of motion properties that our best broadly empirical theories of the world tell us there are. And in this case, it would seem that all of our thought and talk about motion must be systematically false or gappy. Surely, however, this is a conclusion we should want to avoid.

Luckily, we theorists have a way out: we can allow that, following this kind of discovery about the world, we should revise the semantics of motion ascriptions. On the revised view, we allow that contextual supplementation helps our overtly unrelativized motion predicates refer to the available relational motion properties, and thereby allow that ordinary motion ascriptions can be true.

I propose adoption of contextualism about color discourse in the same, revisionary, spirit. Here, too, I claim, our best, broadly empirical empirical account tells us that colors are constituted in terms of relations to (*inter alia*) subjects (as per the argument in §2.1.1). Hence, in order to avoid rendering false or gappy all of the utterances that make use of color predicates predating this broadly empirical discovery, we should accept a revisionary contextualist semantics of color ascriptions (such as that proposed in *RR*) that will allow our overtly unrelativized color predicates to refer to the available relational color properties.²⁸

²⁷This reflects a change in my own view about the meta-narrative of the contextualist proposal of *RR*.

²⁸If I admit (even provisionally) that making linguistic representations of color subject-involving amounts to a revision of linguistic usage, then I cannot hold that data about linguistic usage

2.2.2 Diagnostics for Context Sensitivity

I have admitted that my preferred contextualist explanation of how linguistic color contents could be subject-involving may be descriptively false. Moreover, I claim that this outcome would neither drain the interest of *RR*'s contextualism *qua* revisionary semantics for color predication nor impugn the argumentative purposes to which contextualism is put in that work.

That said, I do not believe that the evidence Allen appeals to in attempting to demonstrate the descriptive inadequacy of contextualism is successful. Let me say why.

Completion: The first diagnostic to which Allen appeals involves completion. The thought here is that expressions whose overt linguistic forms are semantically incomplete, and for which context is thought to supply needed, missing information, typically permit overt specification of the completing material. Thus, for example, while (2a)–(4a) are standardly taken to be semantically incomplete in the relevant sense, each permits supplementation by further linguistic material that results in a the completed counterparts (2b)–(4b).

(2a) The beach is nearby.

(3a) It is raining.

(4a) John's wedding is on Friday at 3.

(2b) The beach is nearby to the cafe.

(3b) It is raining in Seattle.

(4b) John's wedding is on Friday the 23rd of March, 2012, at 3pm PDT.

Likewise, on the contextualism of *RR*, ordinary color ascriptions are semantically incomplete in their overt linguistic form but supplemented by contextually supplied information. Consequently, barring some independent explanation, we should expect that ordinary color ascriptions like (1a) permit expansion into their completed counterparts like (1b). Alas, Allen claims, this prediction fails: (1b) is unacceptable.

The problem with this argument is that, as Allen recognizes (11), linguistic intuitions about the acceptability/unacceptability of (1b) and similar examples are extremely variable. I suggest, then, that completion judgments provide at best weak support for the claim that color ascriptions are context-insensitive.

motivate either (i) a subject-involving construal of perceptual representations of color, or (ii) relationalism about the metaphysics of color properties. But I never attempted to argue for (i) and (ii) in that way. On the contrary, my argument for (i) and (ii) in *RR* comes from considerations about the psychophysics of perception (see §2.1.1). My concern with linguistic usage in *RR* was not driven by a need to argue directly for conclusions (i) and (ii); rather, it was to show how an independently motivated commitment to (i) and (ii) could be squared with facts about the usage of color language.

Context shifting: Allen's second diagnostic for context sensitivity involves the possibility of shifting interpretations of a context-sensitive term within a sentence as a result of shifting the context midway through an utterance of that sentence. Thus, for example, the availability of a true reading of (5) depends on the possibility of interpreting the (context-sensitive) gradeable adjective 'small' relative to two different delineations:

(5) That is a small elephant but not a small animal.²⁹

A second example, which Allen takes from Jackson (2007, 312), involves a sign on the car deck of a ferry containing the instruction:

(6) Cars must not move while the ferry in motion.

Again, it is plausible that this instruction makes sense only because of the possibility of interpreting the two different (context-sensitive) motion predicates it contains ('must not move', 'in motion') relative to two distinct, contextually supplied reference frames.³⁰

In contrast, Allen urges that color predicates do not allow this sort of intrasentential context-shifting behavior with respect to an alleged parameter for perceivers:

For instance, it would seem strange to say:

19) ?The bucket is blue, but the spade is blue-green

if the bucket and the spade are intra-personally indistinguishable in colour to Jack and Jill, but what is meant is that the bucket is blue for Jack, and the spade is blue-green for Jill (13–14).

Once again, I find this argument less than persuasive. As Allen himself notes (14), many paradigmatically context-sensitive expressions — e.g., 'enemy', 'foreigner', modal 'can', and 'nearby' — generally resist this sort of context-shifting as well (cf. Stanley (2005, 72, note 16)). But if we fail to observe the relevant kind of context-shifting with those paradigmatically context-sensitive expressions, the observation that color predicates also typically fail to exhibit context-shifting cannot justify the conclusion that the latter are context-insensitive.³¹

²⁹The delineations relevant to the interpretation of 'small' are plausibly provided (somehow) by the sortals occurring explicitly in (5); I grant, *arguendo*, that this kind of lexical supplementation counts as contextually supplied.

³⁰Again, I'm prepared to concede for the sake of argument that the intended interpretation of (6) depends on an intrasentential context-shift.

³¹I suspect that at least part of what drives linguistic judgments about context-shifting cases (about color predicates and other putatively context-sensitive expressions) is that it can be difficult to pull off contextual shifts in the midst of utterances without compromising felicity. Support for this conjecture comes from the observation that when we move to situations in which the relevant aspect of context — here, contextually relevant perceiver type — are more naturally expected to shift, context-shifting sentences go down more smoothly. Thus, in the vision lab, a normal

Disagreement: Allen’s third and final diagnostic for context-sensitivity involves the possibility of disagreement about color. Briefly, the idea is that, if an expression e is context-sensitive, and if a speaker S_1 applies that expression to a in context C_1 while a second speaker S_2 forbears that expression to a in context C_2 , then it is not guaranteed that there is any single property about whose extension the first and second speakers are disagreeing. For, given the context-sensitivity of e , it could be that S_1 ’s use of e in C_1 and S_2 ’s use of e in C_2 pick out different properties. Thus, the collective indirect disagreement report in (8c) fails when the original utterances in (8a) and (8b) include the context-sensitive ‘nearby’, but (barring ambiguity, irony, etc.) succeeds when we substitute the context-insensitive expression ‘square’:³²

(8a) S_1 in C_1 : a is nearby/square.

(8b) S_2 in C_2 : a is not nearby/square.

(8c) S_3 in C_3 : S_1 and S_2 disagree about whether a is #nearby/square.

Allen takes this not to show that there cannot be collection of the relevant sort involving utterances containing ‘nearby’ or other context-sensitive expressions. Instead, he takes it to show that if there is such collection, then the value of the contextual parameter(s) to which those expressions are sensitive must be shared between contexts C_1 and C_2 (and perhaps also C_3).³³

Of course, disagreements about color (of the sort that can be reported by such collective indirect reports) are rampant — thus, pairs like (9a) and (9b) are common, and appear to legitimate reports such as (9c):

(9a) S_1 in C_1 : a is red.

(9b) S_2 in C_2 : a is not red.

(9c) S_3 in C_3 : S_1 and S_2 disagree about whether a is red.

trichromatic experimenter preparing a stimulus for her protanopic subjects can remark felicitously to another experimenter,

(7) This red patch is grey.

(What Allen says about similar examples (14) suggests that he would reject this case, insisting, with Pritchard (1909), that (7) is elliptical for

(7’) This red patch looks grey.

I just don’t see any reason to hold that this must be so.)

³²Authors who rely heavily on (versions of) this diagnostic for context-sensitivity include (Cappelen and Lepore, 2006; Cappelen and Hawthorne, 2009).

³³Caponigro and Cohen (2011) argue that, strictly speaking, this conclusion is unjustified: for there can be collection over utterances of context-sensitive expressions even when the values of the relevant contextual parameters mismatch. Consequently, we argue, the collection test is unsuccessful as a diagnostic for context-sensitivity in the way that Allen (following Cappelen and Lepore (2006); Cappelen and Hawthorne (2009)) employs it. I’ll put this aside, however, since I doubt that the possibility of disagreements about the application of color predicates depends entirely on the kinds of counterexamples raised by Caponigro and Cohen (2011).

By reasoning parallel to that concerning (8a)–(8c), Allen suggests that the inference involving (9a)–(9c) should be possible, given the contextualist’s view that color predicates are context-sensitive, only if the value of the putative perceiver-type parameter is shared between contexts C_1 and C_2 . But since, he worries, there doesn’t appear to be any reason that that condition will be met, contextualism about color predicates threatens to make ordinary linguistic disagreement about color rare or impossible.

Once again, I remain unpersuaded.

On the contextualist semantics outlined above (§2.2), what is expressed by an utterance in K of a color predicate contains a parameter filled by the K -relevant perceiver-type. Therefore, the perceiver-type parameters relevant to the evaluation of (9a) and (9b) will have the same value only if the C_1 -relevant perceiver-type = the C_2 -relevant perceiver-type. Crucially, however, it is plausible that, in many ordinary contexts, what counts as the contextually-relevant perceiver-type is fairly unspecific and broad — broad enough, in any case, that it will overlap with the contextually relevant perceiver-type of many other ordinary contexts.

None of this is to deny Allen’s contention that there exist pairs of contexts relative to which the contextually-relevant perceiver-types will not match.³⁴ Indeed, I stand by the contextualist prediction that the inference from (9a) and (9b) to (9c) — and, indeed, the intuition that (9a) and (9b) express intercontextual disagreement — breaks down relative to such pairs of contexts. That prediction would be objectionable if it entailed that such inferences always fail relative to an arbitrary pair of contexts, i.e., that there is no intercontextual disagreement about color. But it does not.³⁵

References

- Allen, K. (2012). Colour, contextualism, and self-locating content. *Croatian Journal of Philosophy*.
- Brainard, D. H., Delahunt, P. B., Freeman, W. T., Kraft, J. M., and Xiao, B. (2006). Bayesian model of human color constancy. *Journal of Vision*, **6**, 1267–1281.
- Buchsbaum, G. A. (1980). A spatial processor model for object colour perception. *Journal of the Franklin Institute*, **310**, 1–26.
- Byrne, A. (2010). Experience and content. *Philosophical Quarterly*.
- Byrne, A. and Hilbert, D. R., editors (1997). *Readings on Color, Volume 1: The Philosophy of Color*. MIT Press, Cambridge, Massachusetts.

³⁴Indeed, as Allen notes, I allow (RR, 117–121) that mechanisms of Lewisean accommodation can sometimes make contextually relevant perceiver-types extremely fine-grained. Relative to such contexts, my account predicts that the contextually relevant perceiver-types are less likely to match, hence that an inference parallel to that from (9a) and (9b) to (9c) is correspondingly less likely to go through.

³⁵Thanks to Keith Allen, Dave Barner, Ivano Caponigro, Joshua Gert, Larry Hardin, Katie Wagner, Mike Webster, and Chris Wüthrich for helpful discussion of this material.

- Byrne, A. and Hilbert, D. R. (2007). Truest blue. *Analysis*, **67**(293), 87–92.
- Caponigro, I. and Cohen, J. (2011). On collection and covert variables. *Analysis*, **71**(3), 478–488.
- Cappelen, H. and Hawthorne, J. (2009). *Relativism and Monadic Truth*. Oxford University Press, New York.
- Cappelen, H. and Lepore, E. (2006). Replies. *Philosophy and Phenomenological Research*, **73**(2), 469–492.
- Chisholm, R. (1957). *Perceiving: A Philosophical Study*. Cornell University Press, Ithaca.
- Cohen, J. (2007). A relationalist’s guide to error about color perception. *Noûs*, **41**(2), 335–353.
- Cohen, J. (2012). Perceptual constancy. In M. Matthen, editor, *Oxford Handbook of Philosophy of Perception*. Oxford University Press, Oxford. forthcoming.
- Cohen, J. (2013). Synesthetic perception as continuous with ordinary perception, or: We’re all synesthetes now. In O. Deroy, editor, *Sensory blending: New essays on synaesthesia*. Oxford University Press, Oxford. forthcoming.
- Cohen, J. and Nichols, S. (2010). Colors, color relationalism, and the deliverances of introspection. *Analysis*, **70**(2), 218–228.
- Cohen, J., Hardin, C. L., and McLaughlin, B. P. (2006a). True colors. *Analysis*, **66**(4), 335–340.
- Cohen, J., Hardin, C. L., and McLaughlin, B. P. (2006b). The truth about “The Truth about True Blue”. *Analysis*, **67**(2), 162–166.
- D’Zmura, M. (1992). Color constancy: surface color from changing illumination. *Journal of the Optical Society of America A*, **9**, 490–493.
- D’Zmura, M. and Lennie, P. (1986). Mechanisms of color constancy. *Journal of the Optical Society of America A*, **3**, 1662–1672.
- Egan, A. (2012). Comments on Jonathan Cohen’s *The Red and the Real*. *Analytic Philosophy*.
- Funt, B., Drew, M., and Ho, J. (1991). Color constancy from mutual reflection. *International Journal of Computer Vision*, **6**, 5–24.
- Gilchrist, A. L., Kossifydis, C., Bonato, F., Agostini, T., Cataliotti, J., Li, X., Spehar, B., Annan, V., and Economou, E. (1999). An anchoring theory of lightness perception. *Psychological Review*, **106**(4), 795–834.
- Golz, J. and MacLeod, D. I. A. (2002). Influence of scene statistics on colour constancy. *Nature*, **415**, 637–640.

- Hardin, C. L. (1988). *Color for Philosophers: Unweaving the Rainbow*. Hackett, Indianapolis.
- Jackson, B. (2007). Truth vs. pretense in discourse about motion (or, why the sun really does rise). *Noûs*, **41**, 298–317.
- Jackson, F. (1977). *Perception: A Representative Theory*. Cambridge University Press, New York.
- Johnston, M. (1992). How to speak of the colors. *Philosophical Studies*, **68**, 221–263. Reprinted in Byrne and Hilbert (1997), 137–176.
- Kaplan, D. (2011). Words on words. *The Journal of Philosophy*, **CVIII**(9), 504–529.
- Kennedy, C. (2007). Vagueness and grammar: The semantics of relative and absolute gradable adjectives. *Linguistics and Philosophy*, **30**(1), 1–45.
- Klein, E. (1991). Comparatives. In A. von Stechow and D. Wunderlich, editors, *Semantics: An International Handbook of Contemporary Research*, pages 673–691. Walter de Gruyter, Berlin.
- Kuehni, R. G. (2004). Variability in unique hue selection: A surprising phenomenon. *Color Research and Application*, **29**, 158–162.
- Land, E. H. and McCann, J. J. (1971). Lightness and retinex theory. *Journal of the Optical Society of America*, **61**.
- Lee, H. C. (1986). Method for computing the scene-illuminant chromaticity from specular highlights. *Journal of the Optical Society of America A*, **3**, 1694–1699.
- Lindsey, D. T. and Brown, A. M. (2009). World color survey color naming reveals universal motifs and their within-language diversity. *Proceedings of the National Academy of Sciences of the USA*, **106**(47), 19785–19790.
- MacLeod, D. I. A. (2003). Colour discrimination, colour constancy, and natural scene statistics. In J. D. Mollon, J. Pokorny, and K. Knoblauch, editors, *Normal and Defective Colour Vision*. Oxford University Press, London.
- Malkoc, G., Kay, P., and Webster, M. A. (2005). Variations in normal color vision. IV. Binary hues and hue scaling. *Journal of the Optical Society of America, A*, **22**(10), 2154–2168.
- McGinn, C. (1983). *The Subjective View: Secondary Qualities and Indexical Thoughts*. Oxford University Press, Oxford.
- Pautz, A. (2010). An argument for the intentional view of visual experience. In B. Nanay, editor, *Perceiving the World*. Oxford University Press, New York.
- Peacocke, C. (1984). Colour concepts and colour experiences. *Synthese*, **58**(3), 365–381. Reprinted in Rosenthal (1991), 408–416.

- Pritchard, H. A. (1909). *Kant's Theory of Knowledge*. Clarendon Press, Oxford.
- Rosenthal, D. (1991). *The Nature of Mind*. Oxford University Press, New York.
- Siegel, S. (2010). The weak content view. In B. Nanay, editor, *Perceiving the World*. Oxford University Press, New York.
- Stanley, J. (2005). Semantics in context. In G. Preyer and G. Peter, editors, *Contextualism and Philosophy: Knowledge, Meaning, and Truth*, pages 221–253. Oxford University Press, Oxford.
- Travis, C. (2004). The silence of the senses. *Mind*, **113**(449), 57–94.
- Tye, M. (2006a). The puzzle of true blue. *Analysis*, **66**, 173–178.
- Tye, M. (2006b). The truth about true blue. *Analysis*, **66**(4), 340–344.
- Tye, M. (2007). True blue redux. *Analysis*, **67**(293), 92–93.
- Tye, M. (2012). Cohen on color relationalism. *Analytic Philosophy*.
- Webster, M. A. and Kay, P. (2006). Individual and population differences in focal colors. In R. E. Maclaury, G. V. Paramei, and D. Dedrick, editors, *The anthropology of color*, pages 29–53. John Benjamins Publishing Company.