

Ecumenicism, comparability, and color, or: How to have your cake and eat it, too

Jonathan Cohen*

Couples are wholes and not wholes, what agrees disagrees, the concordant is discordant. From all things one and from one all things.
— Heraclitus (DK B10).

Abstract

Data about perceptual variation motivate the ecumenicist view that distinct color representations are mutually compatible. On the other hand, data about agreement and disagreement motivate making distinct color representations mutually incompatible. *Prima facie*, these desiderata appear to conflict. I'll lay out and assess two strategies for managing the conflict — color relationalism and the self-locating property theory of color — with the aim of deciding how best to have your cake and eat it, too.

Among other considerations, there appear to be two ranges of data that constrain accounts of the ontology of color properties, and that, curiously, seem to pull in opposite directions. On the one hand, considerations about perceptual variation motivate the claim that ascriptions of distinct colors to a single object can be compatibly true, hence, that one object can compatibly exemplify multiple colors (all over and simultaneously). On this sort of view, distinct color ascriptions are compatibly true because the color properties they ascribe are (in some sense to be explained) incomparable, and not mutually constraining. On the other hand, data about linguistic and non-linguistic agreement and disagreement about color motivate the idea that distinct color ascriptions can be comparable and mutually constraining (indeed, incompatible). On its face, this seems like a serious clash. What to do?

Of course, some theorists will reject the characterization of one or both the ranges of data underpinning the clash, and others will accept both ranges of data but will deny one or both of the apparently conflicting theoretical implications drawn from them. This paper will not assess those avenues of response (but see Cohen (2009) for discussion).

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

Instead, it will compare two theories that accept the apparently conflicting demands on our color ontology, thereby attempting to have it both ways.

I'll begin by reviewing the two classes of data, and indicating (briefly) how they motivate the apparently conflicting demands of ecumenicism and comparability (§1). Next, I'll present the two theories — color relationalism (§2) and the self-locating property view (§3) — that are designed to accommodate both desiderata, and that will be our focus in what follows. Because these views can seem very similar, I'll take up the question whether they are merely notational variants, and argue that they are not (§4). Then I'll turn to assessment. I'll bring out some worries for the self-locating view (§5), and I'll defend color relationalism (§6), arguing that the most pressing objections facing the view can be answered. Finally, I'll conclude (§7).

1 Conflicting desiderata

1.1 Perceptual variation and ecumenicism

We can begin by considering data about perceptual variation.¹

Though there's a wide range of data that fits under this heading, the headline description is that there is significant variation in the ways that perceptual systems respond to one and the same color stimulus. It is perhaps easiest to appreciate this point by consideration of a single example of such variation, such as that involving normal visual responses, within a single perceiver, to the central patches in figure 1. The central patches in the two halves of this figure are qualitatively identical in their




Figure 1: The two center gray squares are qualitatively identical in their non-relational properties, but the one against the lighter background appears darker than the one against the darker background.

non-relational properties. However, most subjects report that the central patch looks

¹Though I've presented fuller versions of this argument in several places (including especially Cohen 2009), I confine myself here to a more circumscribed presentation designed only to motivate the ecumenicism desideratum and to get the discussion that follows off the ground.

darker when viewed against the lighter surround in the right hand side of figure 1, and lighter when viewed against the darker surround in the left hand side of figure 1.

This is, of course, but one instance of a very wide-ranging set of intrapersonal variations. A single organism's perceptual responses to a color stimulus vary not only as a function of the surround, but also as a function of the state of adaptation of its retinal receptors, the color temperature, diffusion, and direction of the illumination, the viewing angle, viewing distance, and on and on.

Unsurprisingly given the differences between the visual systems present in different species, there is also significant variation in the perceptual responses of organisms of different species to one and the same stimulus. The chromatic effects that a single stimulus has on these perceivers vary widely as a function of many parameters of their visual systems — retinal cone type populations, population ratios, cone tuning curves, macular and lens pigmentation, and on and on.²

Focusing now just on human perceivers, we can also note that there is evidence of significant interpersonal variation in color perception (even when we put aside perceivers with various anomalies of color vision, and restricting ourselves to normal trichromats). Distinct trichromatic human color perceivers, all of whom pass standard psychophysical tests for normal color vision, have psychophysically distinguishable reactions to one and the same color stimulus.³ On reflection, it is perhaps unsurprising that there should be significant interpersonal differences of this sort. After all, perceptual responses to a color stimulus vary as a function of retinal cone type populations, population ratios, cone tuning curves, macular and lens pigmentation, and on and on.

It would appear, then, that cases of variation with respect to color are ubiquitous and far-ranging. In each such case, whether interspecies, interpersonal, or intrapersonal, one stimulus brings about multiple psychophysically distinguishable effects in perceptual systems. And, given standard (though not universally accepted) assumptions, these multiple, psychophysically distinguishable effects in perceptual systems are states that represent the stimulus.⁴ But if that is so, we can ask: which, if any of the (psychophysically distinguishable) representations of the stimulus is veridical?⁵ The data about perceptual variation motivate ecumenicism because the latter offers the most plausible, general answer to that question.

²For discussion of key findings about interspecies variations in color vision, see, e.g., Jacobs (1981); for some attempts to draw philosophical conclusions from such findings, see Matthen (1999, 2005); Thompson (1995); Thompson et al. (1992).

³One much-discussed instance of this phenomenon is the observed interpersonal variation in spectral loci for the four "unique," or "phenomenally uncomposed," hues of green, blue, yellow, and red. Thus, unique green is that green hue that looks not at all yellowish and not at all bluish; unique blue is that blue hue that looks not at all reddish and not at all greenish; unique red is that red hue that looks not at all bluish and not at all yellowish, and unique yellow is that yellow hue that looks not at all greenish and not at all reddish. Typically, subjects' choices of unique hue loci are intrapersonally remarkably stable (though they may shift over many years, as age changes the filtering properties of the lens and macula), but there is significant variation in the settings made by different (color normal) human perceivers. For general discussion of unique hues, see Hurvich (1981, 66ff); for a good overall review of the evidence of the significant interpersonal variation in the spectral loci for unique hues, see Kuehni (2004).

⁴These standard assumptions are rejected by Smith (2002); Travis (2004), and defended by, e.g., Byrne (2009); Pautz (2010); Siegel (2010a,b).

⁵The question here is metaphysical, not epistemic: it is not 'how do we know which of the perceptual effects veridically represents the stimulus's color?', but 'what makes it the case that one of the perceptual effects (as opposed to others) veridically represents the stimulus's color?'

To see this, consider the question with respect to a case in which there are only two possible variants, such as those arising intrapersonally in the perception of the central patch in the two halves of figure 1. The logically possible answers to the question of which variant is veridical are these: neither, the first to the exclusion of the second, the second to the exclusion of the first, or both. But the first option (neither right) is unacceptably skeptical and revisionary; after all, since we can stipulate that the conditions for color perception are as truth-conducive as we like, this option pushes in the direction of saying that no perceiver in any viewing condition veridically represents the colors of objects. Further, the second and third answers (one representation exclusively) are unacceptably *ad hoc* and arbitrary: since every physical or psychological fact about either variant can be matched with a corresponding physical or psychological fact about the other, it is hard to imagine a well-motivated, principled, and non-question begging reason to believe that either representation is uniquely veridical.⁶ I assume that rational enquirers should avoid both revisionary skepticism and *ad hoc* stipulation when possible. Therefore, these considerations motivate the verdict that, *ceteris paribus*, we should prefer the ecumenical view that both representations are veridical. (And now, of course, we can see that the argument generalizes immediately to cases where there are more than two variants.)

But accepting that, in general, such multiple ascriptions of color to the very same individual are (simultaneously) veridical means accepting that the truth of pairs of these ascriptions is not mutually constraining. If your ascription of *unique green* to *a* and my ascription of *bluish green* to *a* can both be true, that means that your ascription and mine are not incompatible, as they might have seemed. On the contrary, it seems that our ascriptions are incomparable: your ascription attributes a property from a range specific to you, while my ascription attributes a property from a different range specific to me. Of course, we have not yet said just how this incomparability is to be cashed out by a fuller metaphysical account of colors and color ascriptions (after all, we are still at the level of describing and motivating desiderata). But I think the point should be clear enough: the facts of perceptual variation, together with the application of standard rational norms that incline us against either undue skepticism or unmotivated stipulation, motivate the idea that multiple color ascribing variants are non-competitive, incomparable representations of an object's color.

1.2 Agreement and comparability

Of course, data about perceptual variation are not the only empirical considerations that constrain our color ontology.

⁶There may be principled grounds for saying that some of the perceptual variants occurring in some of the cases described above represent the stimulus color erroneously: perhaps, for example, one of the variants arises in a condition that we would want to characterize (for independent reasons) as pathological, hence erroneous — perhaps after the perceiver ingested LSD, or in a perceiver who has undergone blunt instrument trauma to visual areas of her brain, or in a circumstance that is in some important way not ecologically valid. However, it looks as if there will remain significant variation even after we have appealed to all the available principled grounds to exclude as many variants as we can. For example, it is hard to see that there's any independently well-motivated characterization of either of the two perceptual conditions under which we view the central patch in figure 1 that would license setting aside as erroneous the perceptual variant arising under that condition. And given that perceptual variation remains even after we have done all the motivated setting aside of variants possible, it would seem objectionably *ad hoc* to treat the remaining variants that cannot be set aside as systematically misrepresenting the colors of objects.

In particular, one further range of data that we should take into account concerns apparent agreement and disagreement in our linguistic exchanges about color, as might occur in the following sort of perfectly ordinary (if boring) discourse:

- (1) a. S_1 : a is unique green.
- b. S_2 : a is not unique green (/ a is bluish green).
- c. S_3 : a is unique green.

The overwhelmingly natural description of this discourse is that, in uttering (1a–c), S_1 and S_2 are *disagreeing* with one another, and S_1 and S_3 are *agreeing* with one another. But, at least on standard views, S_1 and S_2 can't succeed in disagreeing unless the content of S_1 's ascription in (1a) *conflicts with* that of S_2 's ascription in (1b); and, similarly, S_1 and S_3 can't succeed in agreeing unless the content of S_1 's ascription in (1a) is *identical to* the content of S_3 's ascription in (1c).⁷ All of this is to say that, at least on standard views, considerations about interpersonal agreement and disagreement require that different speakers/perceivers can endorse or reject the very same content.

Intrapersonal considerations about agreement point in the same direction as well. Thus, it is plausible that without the capacity to represent the very same color property twice, a single thinker would be unable to reason without equivocation through premises that mention colors. And, indeed, there are parallel motivations that don't even require the assumption that the organism engages in anything as sophisticated as reasoning. Thus, for example, it's hard to see how to explain perceptual object recognition (which appears to occur throughout the animal kingdom (Jitsumori and Delius 2001; Soto and Wasserman 2010; Spetch and Friedman 2006)) without the idea that an organism can ascribe to perceived objects on multiple occasions the very same (*inter alia* color) properties (cf. Byrne and Hilbert 2003, p. 58). Similarly, it's hard to make sense of classical conditioning without the idea that organisms can represent that the currently perceived object exemplifies the very same properties (e.g., color properties) again, or else exemplifies properties that are incompatible with those perceived earlier.

All told, then, these considerations suggest a lesson that is directly at odds with the conclusion derived above from the facts of perceptual variation. Namely, they suggest that distinct color ascriptions *are* mutually constraining, or comparable, after all.

1.3 Clash

We have, then, empirically well motivated reasons to accept both of what are apparently clashing desiderata on theories of color. On the one hand, it would seem that considerations about interpersonal and intrapersonal (and interspecies) perceptual variation motivate the idea that distinct representations of color should be incomparable, or not mutually constraining. And, on the other, it would seem that considerations about interpersonal and intrapersonal (and interspecies) agreement motivate the idea that distinct representations of color should be comparable, or mutually constraining. As usual, we should aim to formulate theories of our target (in this case the nature of color) that satisfy known, empirically well-motivated, desiderata. But the present situation seems more than usually problematic: what should we do in cases, like this one, where the desiderata appear to come into direct conflict?

⁷For reasons for believing that this standard view about agreement/disagreement may be too simple, see Caponigro and Cohen (2011). I'll ignore such complications here, since the confounding factors discussed by Caponigro and Cohen aren't at issue in the cases under consideration.

Many — perhaps most — writers who have considered these matters have responded to the clash by rejecting one or both of the desiderata described in §1. (From my unsystematic observations, philosophers seem far more eager to give up on ecumenicism than on comparability.) For what it is worth, my own view is that those sorts of response are unpromising.⁸ But that’s not what this paper is about.

This paper is about views that accept the apparently conflicting desiderata at face value and attempt to respect both of them. In particular, in what follows I will present and then compare two (and a half — see note 24) competing color ontologies that accept both desiderata, and so attempt both to have their cake and eat it, too.

2 Color relationalism

One view that attempts to respect both desiderata is the color relationalist account I have defended elsewhere (Cohen 2009).

The leading idea behind relationalism is that colors are constituted in terms of relations to perceiving subjects and perceptual circumstances. For the relationalist, colors are not, fundamentally, monadic, intrinsic, properties of objects. Rather, on this view, an individual particular has a color by having the right sort of relation (whatever that amounts to) to perceivers and circumstances — for example, an object a might exemplify the color *red* for S_1 in C_1 by virtue of bearing the right relation (whatever that amounts to) to the perceiving subject S_1 in the perceptual circumstance C_1 .⁹

It may be helpful to compare the relationalist’s way of thinking about colors with properties that are less controversially relational, such as *being a sister*. One of the things it means to say that this property is relational is to say that nothing is a sister simpliciter: something is a sister of *Beatrice*, or a sister of *Chris*. And things that are sisters of Beatrice or Chris get to be that way by virtue of bearing the right sort of relation (whatever that amounts to) to Beatrice or Chris. So, too, claims the color relationalist, for color properties.

2.1 Color relationalism and ecumenicism

Color relationalism is well positioned to respect the ecumenicism desideratum. The crucial observation here is that if an object bears a relation R to a sequence of relata $\langle r_1, r_2, \dots \rangle$, that entails nothing at all about whether it does or does not bear R to some other sequence of relata $\langle r'_1, r'_2, \dots \rangle$, or whether it does or does not bear some different relation R^* to either one of those sequences. Thus, if Alice is a sister of Beatrice, this neither requires nor precludes that Alice is a sister of Chris. (Nor, indeed, does it require or preclude that Alice bears some other relation to either Beatrice or Chris). So, too, for color relationalists, if a is unique green to S_1 in C_1 , this neither requires nor precludes that a is unique green to S_2 in C_2 . (Nor, indeed, does it require or preclude that a bears some other color constituted in terms of a relation to either S_1 or S_2 .) For relationalists, then, objects can compatibly exemplify many distinct colors (all over and

⁸See Cohen (2009, ch. 3) for critical consideration of several of the most important responses of this kind.

⁹ Here I characterize relationalism in a way that is agnostic about the nature of the color constitutive relation, since I think it is useful to consider the broader relationalist framework while allowing that different relationalists might disagree about the best way of filling in the details. Because color relationalism is formulated in a way that leaves this room for debate, it is probably best regarded as a view-family rather than a single determinate view.

at the same time) for the simple reason that they can compatibly exemplify multiple relations to distinct perceivers and perceptual circumstances. It is no wonder, then, that relationalists can accept the truth of multiple ascriptions of color to one and the same object. For relationalists, multiple ascriptions of color to an individual are incomparable in just the sense desired: they are not mutually constraining.¹⁰

2.2 Color relationalism and comparability

But if relationalism is well-suited to account for the ecumenicism desideratum, it appears — at least in the form presented so far — singularly ill-suited to account for the comparability constraint. As just noted, relationalists secure the compatibility of multiple color ascriptions by treating these as ascriptions of properties constituted in terms of relations to subjects and viewing conditions. But if that is so, then it would seem that relationalists will be forced to hold that the properties ascribed to *a* by the three utterances in (1a–c) will, likewise, be constituted in terms of relations to the perceivers who utter them in the circumstances they are in while making the utterances. I.e., this suggests that the properties ascribed by (1a–c) are those listed in (2a–c) (respectively):

- (2) a. *unique green to S₁ in C₁,*
- b. *not unique green to S₂ in C₂/bluish green to S₂ in C₂,*
- c. *unique green to S₃ in C₃.*

And now the trouble for relationalism is immediately apparent. The relational properties listed in (2a–c) are incomparable: *a*'s exemplification of *unique green to S₁ in C₁* doesn't conflict with or otherwise constrain *a*'s exemplification of *not unique green to S₂ in C₂/bluish green to S₂ in C₂*, nor does it agree with or otherwise constrain *a*'s exemplification of *unique green to S₃ in C₃*. It would seem, then, that relationalism is egregiously unable to meet the comparability desideratum.

The problems relationalism has with comparability go to the heart of the view. Indeed, they are nothing more than the flip side of the view's successes with perceptual variation. Relationalism is positively motivated by a desire to preclude comparability in the contents of color ascriptions so as to avoid the need for stipulative choices in cases of perceptual variation. It is, therefore, no surprise that, having responded to those problems by repudiating content comparability, relationalists have a hard time securing comparability in other cases where it is wanted (e.g., to account for the apparent agreement and disagreement we see in (1a–c)).¹¹ For these reasons, it seems to me that

¹⁰The schematic description in the main text invites the question: what parameters individuate perceivers and perceptual circumstances?

Again, I can imagine different relationalists answering that question differently. However, in so far as the position aims to accommodate the kind of ecumenicism motivated by the facts of perceptual variation, there is a reason to think we should include any parameter variation along which affects the psychophysical effect of the stimulus in the perceiver and that cannot be excluded in a principled and theory-independent way. Following this procedure will plausibly require that we take account of parameters including properties of the chromatic/achromatic surround, properties of the illumination, viewing size and distance, simultaneously seen objects, retinal cone type populations and ratios, state of adaptation of the visual system, and so on. Though this is clearly an empirical matter, it is reasonable to expect that this strategy will result in descriptions of perceivers and perceptual conditions, hence of color properties, that are significantly more fine-grained than we would have come up with prior to investigation.

¹¹Tye (2012) offers this consideration (among others) as a reason to abandon the color relationalism of Cohen (2009); however, he does not comment on any of the strategies for

relationalism, if left in the sparse form in which it has been stated so far, has no hope of meeting the comparability desideratum.

Luckily, we can salvage the view by supplementing it.

The supplementation I suggest is the addition of a further level of color-representation — a level additional to the content-incomparable perceptual representation of fine-grained relational properties that we have already discussed. Namely, the relationalist should allow, additionally, that perceivers' cognitive systems (and linguistic systems) represent what Cohen (2009) calls 'coarse-grained colors' — relational properties analogous to the fine-grained colors discussed above, but whose parametric positions are filled by relatively unspecific, or less determinate, specifications of perceivers and perceptual conditions.

The addition of this new, coarse-grained, level of color content may initially seem unparsimonious, but it is something that relationalists have independent reason for accepting. After all, ordinary color predications in color language lack overt parameters for the perceiver/circumstance relata that relationalists take to partially constitute colors; and it would implausibly over-intellectualize our lexical/semantic resources (which presumably predate sophisticated color science) to suppose that these linguistic representations harbor hidden variables for the many fine-grained parameters that appear to be perceptually relevant. So it makes sense for relationalists to allow that ordinary cognitive and linguistic representations of color have contents that are less determinate than the fine-grained colors.¹² On the contextualist version of this strategy I favor, the specifications of perceivers and perceptual circumstances are contextually supplied supplementations to the much sparser information expressed explicitly in language or thought. Thus, on this view, the predicate 'is yellow', as uttered/thought in context K , expresses the property *yellow for the perceivers relevant in context K under the perceptual circumstances relevant in context K* .¹³ ¹⁴ And similarly for other color predicates.¹⁵

responding to this criticism, including those discussed in that work (and below), so it is not clear whether or why he thinks those strategies are inadequate. For further discussion of Tye's objections, see Cohen (2012).

¹²Besides providing a plausible account of cognitive/linguistic representation of colors, and additionally allowing for a solution to the initial clash between desiderata that is my focus in this paper, the introduction of coarse-grained colors also paves the way for relationalist responses to an array of otherwise troublesome objections against relationalism having their source in our ordinary thought and talk about color — e.g., the worry that relationalism legitimates more color attributions than we would ordinarily accept, that it is overly permissive in the color attributions it licenses, and that it precludes errors of color representation. For details, see Cohen (2009, chapter 4).

¹³Though this wasn't explicit in Cohen (2009), I have come to think that this contextualist semantics is best understood as a self-consciously revisionary proposal about how to hook overtly unrelativized color predicates onto the world, given the ontological inventory color relationalism is committed to (for reasons motivated by perceptual rather than linguistic phenomena).

¹⁴In cases where a perceiver S_1 in a perceptual circumstance C_1 takes herself to be a K -relevant perceiver and C_1 to be K -relevant circumstance, she may, on the strength of her perceptual representation of a as exemplifying the fine-grained property *yellow to S_1 under C_1* , come to hold a cognitive/linguistic representation of a as being yellow *simpliciter* — viz., as exemplifying *yellow for the perceivers relevant in context K under the perceptual circumstances relevant in context K* . Whether this transition between representations is epistemically warranted will depend on, possibly among other things, whether S_1/C_1 are indeed K -relevant, as S_1 takes them to be.

¹⁵What types of perceivers/circumstances are relevant in a context K ? If K is a more or less ordinary conversational context in which there are no special presuppositions in force, it is plausible that the relevant perceivers/circumstances are something like the (metaphysically

Though this story has a number of significant advantages, the one most germane for our purposes is that, because the parametric positions for perceivers and perceptual conditions in such coarse-grained colors are filled in relatively unspecific (and context-relative) ways, ascriptions of such properties will exhibit greater comparability than do representations of fine-grained colors. To see this, consider the discourse in (1a–c) once again. Appealing to the contextualist semantics just outlined, we can hold that these ascriptions, if made in a common context K , ascribe to a the following properties (respectively):

- (3) a. *unique green to a K -relevant perceiver in K -relevant perceptual circumstances,*
 b. *not unique green to a K -relevant perceiver in K -relevant perceptual circumstances*
(bluish green to a K -relevant perceiver in K -relevant perceptual circumstances),
 c. *unique green to a K -relevant perceiver in K -relevant perceptual circumstances.*

Crucially, however, and unlike the incomparable properties in (2a–c), the properties in (3a–c) are straightforwardly comparable (despite being relational): property (3a) conflicts with/is impossible with property (3b), and is identical with property (3c).

The upshot, then, is that color relationalism, once supplemented in independently motivated ways, contains resources to meet both the ecumenicism and comparability desiderata. On the one hand, relationalists provide for ecumenicism by pointing to the incomparability of perceptual representations of fine-grained colors. On the other, they can account for comparability by appeal to their separate, cognitive/linguistic level of coarse-grained representation. By postulating both forms of color representation, and by distributing the explanatory labor between these resources, relationalists can answer to both desiderata, despite the apparent clash between them.

3 The self-locating property theory of color

Egan (2006, 2010, 2012) and Brogaard (2012, 2014) offer an interesting alternative color ontology that (among its other advantages) holds out the prospect of simultaneously accommodating the ecumenicism and comparability desiderata.¹⁶ I'll follow Brogaard

unprincipled) “normal” perceivers/circumstances that traditional secondary quality theorists have invoked — perhaps perceivers more or less similar to most of the conversational participants themselves (perhaps members of the same species, or those who make similar color discriminations most of the time), and circumstances more or less similar to most of the circumstances the conversational participants encounter (actually, nowadays, and hereabouts). But I take it that conversants can, if they wish (by stipulation, presupposition and conversational accommodation, etc.), restrict the range of conversationally relevant perceivers/circumstances in any other way that serves their needs — perhaps to perceivers who are dichromats, deuteranopes, women, pigeons, non-anomalous trichromats adapted to a stimulus used in their psychophysics lab, molecular duplicates of Barack Obama, or what have you, and to circumstances involving a particular viewing angle, adaptation pattern, illuminant, chromatic surround, etc., or any combination of such parameters.

Of course, while the particular coarse-grained colors represented in a context K will ordinarily serve the context-dependent conversational interests of the representers present in K , this doesn't mean that the K -relevant perceivers/perceptual conditions are distinguished from other sorts of perceivers/perceptual conditions in any metaphysically significant way. That one coarse-grained color is represented rather than another does not amount to a metaphysically principled choice of one perceptual variant over others in the context of the argument from perceptual variation, and therefore in no way relieves the pressure to accept fine-grained colors in our ontology.

¹⁶For the record, Egan's endorsement of the view is tentative (cf. Egan 2012, 310, note 1).

in referring to this view as the *self-locating property theory of color*. In this section I'll set out the view, and show how it answers to our desiderata.

3.1 Colors as self-locating properties

First a bit of set up.

We can start with a standard Stalnakerian conception of contents/objects of propositional attitudes as classes of possible worlds (Stalnaker 1979) — usually thought of as the worlds where the content in question is true. Though this standard picture works well for many purposes, it gives out in describing propositional attitudes that are irreducibly *de se*, such as this famous case from Lewis (1979a):

Consider the case of the two gods. They inhabit a certain possible world, and they know exactly which world it is. Therefore they know every proposition that is true at their world. Insofar as knowledge is a propositional attitude, they are omniscient. Still I can imagine them to suffer ignorance: neither one knows which of the two he is. They are not exactly alike. One lives on top of the tallest mountain and throws down manna; the other lives on top of the coldest mountain and throws down thunderbolts. Neither one knows whether he lives on the tallest mountain or on the coldest mountain; nor whether he throws manna or thunderbolts (Lewis 1979a, pp. 520–521; cf. Perry 1979; Peacocke 1979).

The problem for the standard view is that it doesn't seem to make room for something for the two gods to be ignorant of — by stipulation, both know every proposition true at their world. Lewis's solution to this problem was that we should think about objects of attitudes not as classes of worlds, but as classes of *centered worlds* — i.e., classes of world, time, individual triples. This solves the problem of the two gods neatly by giving them something to be ignorant of (despite their propositional omniscience): each god may know that he is in world w , but can be ignorant of whether his world is centered on $\langle w, t, \text{god}_1 \rangle$ or $\langle w, t, \text{god}_2 \rangle$. Lewis and others working in this tradition propose that it won't hurt to think of all content in this way, though non-*de se* contents will turn out to be harmlessly degenerate special cases — what Egan (2012, p. 310) calls “boring centered world propositions” — that don't distinguish between positions within a world, and include either every $\langle w, t, i \rangle$ or no $\langle w, t, i \rangle$ sharing the world component w .

The transition from identifying contents with worlds to centered worlds goes naturally with a parallel transition in the understanding of the properties attributed in perception/belief/etc. Traditionally, we think of properties as mappings from possible worlds to extensions: thus, property F will map a world w to the class of objects exemplifying F in w . But now we can think of properties as mappings from *centered worlds* to extensions: thus, property F will map a centered world $\langle w, t, i \rangle$ to the class of objects exemplifying F in $\langle w, t, i \rangle$. Equivalently, we can treat properties as mapping from objects to centered worlds propositions: thus, property F will map each object a to the class of centered worlds $\langle w, t, i \rangle$ with respect to which a exemplifies F . On the new view, then, objects won't exemplify properties with respect to a world as a whole, but, rather — and as the label ‘self-locating’ is intended to evoke, with respect to a world *qua* centered on an individual and a time. (Once again, there will be degenerate/boring cases where a property has the same extension with respect to every $\langle w, t, i \rangle$ whose world component is w .)

And now we can, at last, state the self-locating property theory of color: this is the view that colors are non-boring centering features. Thus, for example, the property

green is a mapping from any object x to the single centered world proposition true in $\langle w, t, i \rangle$ iff x bears the right relation (whatever that amounts to) to i in the circumstances i occupies at t in w .¹⁷

3.2 Self-locating properties, ecumenicism, and comparability

Among its many virtues, the self-locating property theory of color provides for elegant explanations of both ecumenicism and comparability, and does so without requiring two separate levels of color representation in the way that relationalism does.

We can see this by considering how the self-locating property view would understand the ascriptions in the discourse (1). On this theory, the utterances of the three successive ascriptions (1a–c) deliver the following centered world propositions, respectively:¹⁸

- (4) a. the centered world proposition true in $\langle w, t, i \rangle$ iff a is disposed to look unique green to i in circumstances i occupies at t in w .
- b. the centered world proposition true in $\langle w, t, i \rangle$ iff a is not disposed to look unique green to i in circumstances i occupies at t in w (/the centered world proposition true in $\langle w, t, i \rangle$ iff a is disposed to look bluish green to i in circumstances i occupies at t in w).
- c. the centered world proposition true in $\langle w, t, i \rangle$ iff a is disposed to look unique green to i in circumstances i occupies at t in w .

This treatment delivers comparability in terms of the relations between the conditions appearing on the right hand sides of (4a–c). For example, the view allows us to say that the ascriptions (1a) and (1b) conflict because the condition mentioned in (4a) and that mentioned in (4b) are disjoint: there’s no centered world $\langle w, t, i \rangle$ such that the extensions determined by *unique green* and *not unique green* (/bluish green) relative to $\langle w, t, i \rangle$ overlap. Nothing can satisfy both of these conditions with respect to any $\langle w, t, i \rangle$, because nothing can be both disposed to look unique green to i in circumstances i occupies at t in w and not disposed to look unique green (/disposed to look bluish green) to i in circumstances i occupies at t in w . Correspondingly, the view can explain the sense in which (1a) and (1c) agree in terms of extensional overlap: the condition in (4a) and that in (4c) determined, respectively, by attributions (1a) and (1c), match in their extensions relative to any $\langle w, t, i \rangle$ (indeed, the “two” conditions here are identical).

On the other hand, the view also makes good on the ecumenicism desideratum, allowing for variation without error. This is because, even when the contents of two ascriptions pick out disjoint extensions relative to a given center, as we saw in the case of (1a) and (1b), they can nonetheless pick out different, and not mutually constraining, extensions relative to distinct centers. Though there is no centered world $\langle w, t, i \rangle$ such that the extensions determined by *unique green* and *not unique green* (/bluish green)

¹⁷ The formulation of Egan (2012) commits to the more specific idea that the color constitutive relation is a disposition to look green: “Attributing being green to Kermit delivers the centered worlds proposition that’s true in $\langle w, t, i \rangle$ iff Kermit is disposed to look green to i in the circumstances i occupies at t in w ” (311). (Brogaard (2014) doesn’t commit to any particular version of the view that colors are self-locating.) I have no specific objection to making such further commitments, except to note that, just as I observed in connection with relationalism (note 9), they are separable from the proposal to treat colors as self-locating properties.

¹⁸Here I adopt Egan’s preferred precisification of the view (cf. note 17) to smooth exposition; nothing essential hangs on this choice.

relative to $\langle w, t, i \rangle$ overlap, the properties can compatibly determine extensions relative to distinct centers, and there is no reason these extensions might not overlap. Thus, S_1 's utterance of (1a) is naturally evaluated relative to the centered world $\langle w, t, S_1 \rangle$, and S_2 's utterance of (1b) is naturally evaluated relative to the centered world $\langle w, t, S_2 \rangle$. And while, to repeat, nothing can satisfy the conditions picked out by the two properties relative to one and the same center, an individual a can, compatibly, satisfy each condition relative to a distinct center: a can, compatibly, be both disposed to look unique green to S_1 in the circumstances S_1 occupies at t in w and disposed to look bluish green to S_2 in the circumstances S_2 occupies at t in w . Since, on this view, a can compatibly satisfy the truth-conditions of multiple ascriptions in cases of perceptual variation, there is no need to choose between them. Ecumenicism found.¹⁹

4 Non-collapse

Color relationalism and the self-locating property theory of colors share many features — many of them stemming from the fact that both aspire, as many (most?) other accounts of color do not, to respect both of the desiderata discussed in §1. Indeed, so similar are the views that it is possible to wonder whether they are mere notational variants. I believe they are not.

The most important respect in which the two views look similar is that both make crucial appeal to a relative or relational, subject-involving, element as a way of providing for the ecumenicism desideratum. But the views locate this element in importantly different ways. Relationalism locates the subject-involving element in the nature of color properties: it says that the properties constitutively involve relations to subjects. The self-locating property view, in contrast, effectively treats color properties as non-relational, but claims that the way these non-relational properties pick out extensions is subject-/center-dependent. So where the relationalist treats the relevant subject-dependence by building subject parameters into the nature of the color properties, the self-locating property view treats subject-dependence by building it into the relation between properties and their extensions.

In fact, even though I just said that the self-locating property view effectively treats colors as non-relational, that somewhat overstates the commitments of the view. It is

¹⁹It is worth being clear that the sense of conflict/comparability that the self-locating property view provides, as explicated in the main text, is revisionary. The traditional reading of the conflict intuition with respect to a pair like (1a–b) is that what is expressed by (1a), as uttered by S_1 , is *impossible with, hence rules out* what is expressed by (1b), as uttered by S_2 . Whereas, on the current proposal that S_1 's utterance of (1a) and S_2 's utterance of (1b) express the centered worlds propositions given in (4a–b), what is expressed by the former is *not* impossible with, and does not rule out, what is expressed by the latter. (That is, after all, why the view counts as respecting the ecumenicism desideratum.)

But surely it would be unfair to reject the self-locating property theory simply because it fails to accommodate the conflict/comparability intuition (or the underlying agreement/disagreement data) in the traditional way. The self-locating property theorist's understanding of comparability, given in terms of disjointness/overlap of the conditions appearing in the contents expressed by color predications, is offered as a *replacement* for the traditional, compatibility-based understanding — and, importantly, one that collapses onto the traditional understanding in cases where expressions are evaluated with respect to just a single center, so accounts for the appeal of the traditional notion. Rejecting the proposed inheritor on the grounds that it doesn't exactly match the original it is intended to replace would require imbuing the original intuition with a level of theoretical specificity we have no reason to expect of it, and would amount to begging the question against the replacement.

true that the self-locating property story is compatible with the understanding the color properties as non-relational, and its proponents employ locutions that suggest that that is what they have in mind. But we should notice that the self-location part of the story — the part that is doing the work in securing ecumenicism, and that makes the view seem similar to relationalism — is actually *agnostic* about whether colors are relational or not.

This agnosticism means that the self-locating property view rejects the claim at the very center of color relationalism. Moreover, as I'll discuss below, the structural differences between the two accounts make them vulnerable to very different sorts of objections. Given these differences, it's hard to see the two views as mere notational variants.

5 The self-locating property view reconsidered

The self-locating property view is clearly an important and interesting competitor to color relationalism with many merits, not least of which is that it provides ingenious resources for meeting the desiderata under discussion here. However, the view invites some worries that are worth getting out in the open. In particular, I believe the most important worries about the view are relatively big picture concerns about its motivation and overall explanatory adequacy. To be clear, I do not believe the considerations that follow constitute anything like a decisive refutation of the self-locating property theory: readers will have to decide for themselves just how important the worries are, and just how telling they are against the theory. But I believe they are relevant to any clear-eyed assessment of the account.²⁰

5.1 A bridge too far?

The apparatus of self-locating properties is powerful and elegant, and is extremely useful (perhaps even correct) as an account of some properties.²¹ But I find it not obvious that the view should be extended to color properties, in particular.

One reason for pause in this regard is the observation that the self-locating property view is equally applicable to paradigm cases of relational properties, such as motion properties. In these post-Einsteinian days, we have become accustomed to the idea that motion properties are constituted in terms of relations to reference frames — even if we sometimes fail to specify the latter overtly in our verbal motion ascriptions. Thus, the standard, relationalist story about motion is that *a* does not exemplify *moving at 60mph simpliciter*, but, rather, *moving at 60mph relative to frame F_1* . But notice that we could, if we choose, deny what is usually thought of as a relationalist insight about motion properties, and instead describe the situation in terms of centered worlds. The trick here, just as in the application of the apparatus to colors, is to locate the subject-/frame-/center-dependent element not in the nature of the properties themselves (as per relationalism), but in the way that the properties pick out extensions. Thus, we could say that the ascription (5a) delivers the content (5b)

- (5) a. *a* is moving at 60mph.

²⁰Material in this section expands on remarks in Cohen (2012).

²¹E.g., I find it extremely attractive as an account of the doxastic property that distinguishes the two gods of Lewis (1979a), or of the doxastic property shared by all of the people who believe their own pants are on fire (Kaplan 1989).

- b. the centered world proposition true in $\langle w, t, i \rangle$ iff a is moving at 60mph relative to the frame i occupies at t in w .

So it is possible, using the apparatus of self-locating properties, to describe motion properties in a way that respects the relativity to frames of reference that motivates the special theory of relativity, but that denies their relationality. But this seems to be something of a revisionary characterization of the relevant episode in intellectual history. Instead of conceiving of the episode as involving a new discovery about the (surprising) kinds of properties motion properties are, we are describing it as involving a discovery about how the old, familiar motion properties we knew and correctly understood all along turn out (surprisingly) to pick out their extensions.

Moreover, there is a sense in which the self-locating property description of motion properties is not only revisionary, but serves our epistemic-explanatory needs less well than the standard, relationalist description. For, supposing that hard-won inquiry has revealed the need for a subject-involving element in understanding a certain property (as proponents of relationalism and self-locating property views will agree), it seems odd to prefer a metaphysics on which that property is treated as possibly not subject-involving (but where the way it selects extensions is subject-involving). On the contrary, under these circumstances, it seems more useful to mark the hard-won knowledge by encoding the property's relationship to subjects directly in our metaphysical account of its nature. The self-locating property view is (at worst) guilty of errors of omission here, since it is deliberately noncommittal about the metaphysical natures of the properties to which it applies. My point, however, is that under the imagined circumstances where our best theories of the world (broadly speaking) warrant such commitments, we should prefer a metaphysically committal view over metaphysical quietism.²²

²²This consideration might be thought to cut the other direction as well. Thus, if you accept, with such authors as Dancy 1986, p. 181; Armstrong 1987, p. 36; Boghossian and Velleman 1989, p. 85; Chalmers 2006, p. 556; Gibbard 2006, p. 10; Hazlett and Averill 2010, that phenomenology presents colors as non-relational, then you might take this as a reason to prefer the self-locating property theory on the grounds that it does *not* require locating the subject-involving aspect of colors in the metaphysics of the properties themselves.

I find this consideration unpersuasive.

For one thing, it's not clear that the self-locating view provides any advantage over relationalism in the respect contemplated. For the sort of phenomenological evidence that fails to reveal colors as relational also plausibly fails to reveal colors as self-locating. As such, the phenomenological evidence appears to constitute a *prima facie* threat against both views, rather than providing a reason to prefer either one over the other. Of course, one could reply that this apparent symmetry is misleading. E.g., perhaps one could argue that phenomenology is both committal and revelatory about the metaphysics of the properties it represents — hence that the absence of phenomenal evidence of relationality is evidence of the absence of relationality, while phenomenology is mute about whether the properties it represents are self-locating — hence that the absence of phenomenal evidence of self-location is *not* evidence of the absence of self-location. But those claims about phenomenology are hardly self-evident, and, at a minimum, deserve defense before we can take the phenomenological evidence as favoring the self-locating view over relationalism.

For another, as I have urged elsewhere (Cohen 2009, ch. 6), we should not accept the claim grounding the objection under consideration to the effect that phenomenology fails to reveal colors as relational. It's plausible enough that introspection on isolated, punctate phenomenal episodes (what Levin (2000) calls "glances") fails to disclose evidence of the relationality of colors. But there's no reason to expect phenomenology of that kind to speak to questions about relationality or other elements of the metaphysical makeup of target properties, so our failure to detect evidence of relationality from that sort of phenomenology is in no way decisive. On the other hand, the sort of phenomenal evidence that can reasonably be expected to speak to such questions — evidence that requires phenomenal comparison between multiple experiential episodes, possibly together with

5.2 Quietism and systematization

We noted (§4) that the self-locating theory is officially non-committal about whether colors are relational. But the view's agnosticism about relationality reflects a deeper difference between it and color relationalism. The self-locating property view is, in a way that relationalism is not, trenchantly quietist about the metaphysics of color properties. True, it says that color properties are interesting rather than boring self-locating properties — that they pick out different extensions relative to different centers. However, it says nothing about the nature of color properties that would explain why this is so about them, in particular. Indeed, it is compatible with the truth of the self-locating property view that there should be no explanation, or many different explanations, for this behavior — which has been claimed, after all, for a pretty diverse range of properties.²³

One particular consequence of this quietism is that the self-locating property theory lacks resources to explain systematically why the many different colors are, one and all, subject-involving. I take it that the case for a subject-involving element in one color is not substantially better or worse than the case for a subject-involving element in any other. The relationalist is, at least in principle, poised to explain this commonality by pointing to shared subject-implicating features or structure in her preferred metaphysics of the different color properties (naturally, different sorts of relationalists will draw here on different apparatus). In contrast, because of the self-locating property view's agnosticism about the underlying metaphysics of color properties, there is nothing in that view that explains why different colors all turn out to select extensions in a subject-involving way. Indeed, it is not only that different colors all happen to select extensions in a subject-involving way, but they appear to select extensions in ways that are systematically dependent on the very same sorts of subject-involving parameters — viewing distance, cone populations, state of retinal adaptation, etc. The present point is that, while the self-locating property theorist can model the subject-involving extension selections that color properties make, there is nothing in her metaphysics of color to explain why all of the different colors are appropriately modeled in that very same way.²⁴

some amount of ratiocinative reflection — *does* provide evidence of relationality (as we saw when we appealed to exactly this sort of evidence for just this reason in §1.1).

²³Candidates endorsed by Brogaard and/or Egan would include at least doxastic properties (e.g., the property shared by all those who believe their own pants are on fire), properties picked out by “predicates of personal taste” such as ‘tasty’ or ‘fun’, egocentric spatial properties (*being to the right*, *being such that the tree is further away than the house*), and the indiscriminating property *being self-identical*.

²⁴ Another possibility that suggests itself is that one might construct a hybrid view combining a relationalist metaphysics of color properties with a self-locating property story about the property extensions and semantics.

This sort of a view might start by accepting a fine-grained inventory of relational colors (*red for S_1 in C_1 , green for S_2 in C_2 , etc.*). It could then add unrelativized/non-relational but (interestingly) self-locating colors: we could say that attributing *red* to *a* delivers the centered world proposition true with respect to $\langle w, t, i \rangle$ iff *a* exemplifies the fine-grained relational property *red to i in the circumstance i occupies in w at t* .

This hybrid view would meet the ecumenicism desideratum in two ways, corresponding to the explanations of ecumenicism supplied by both the of the theories it draws upon — viz., because of its reliance on fine-grained relational properties constituted in terms of relations to subjects and circumstances, and also because distinct properties will be compossible relative to distinct centers) It would also meet the comparability desideratum because of its reliance on self-locating properties: it makes sense, on this view, to ask about whether the extensions of two self-locating

6 Relationalism reconsidered

Though there are many worries that can be raised against color relationalism, there are at least three clusters of objections that seem especially pressing against the backdrop of the problems we are discussing, and that therefore merit further discussion here: worries about parsimony, about the adequacy of the relationalist story about color representation, and about especially hard cases concerning comparability. It's worth noting that these worries are less about the big motivational/explanatory picture than those considered with respect to the self-locating property theory in §5, and more about showing that relationalism can avoid what might otherwise seem to be objectionable, but local, consequences. In any case, these more local problems are the ones that seem to me (and, to judge from the literature, to others) the most serious kinds of threats to color relationalism. In this section I'll argue that these problems are less serious than they might appear, and not reasons to give up on the view.

6.1 Parsimony

A first set of reasons for objecting to relationalism arises from concerns about parsimony.

Thus, consider first relationalism's apparently unparsimonious proliferation of distinct relational properties. As we saw in §2, the key to the relationalist's capacity to deliver ecumenicism is her claim that a single object can compatibly (simultaneously, all over) exemplify distinct relational properties. While one can argue about whether this strategy commits the relationalist to attributing to ordinary objects infinitely many, or merely a large finite number of colors (cf. Cohen 2009, pp. 133–134), it's certainly possible to fear that the view is, in this respect, less than perfectly parsimonious in its inventory. More to the point, one might think that the view's way of delivering ecumenicism is less parsimonious than that of the self-locating property account.²⁵ Should we, therefore, not prefer the self-locating property theory over relationalism in that, while both secure ecumenicism, only the former does so without an unlovely proliferation of properties?

color properties, evaluated relative to the very same center, overlap or are disjoint. Moreover, the hybrid view would allow, as the self-locating property view does not, a systematic explanation of *why* its self-locating color properties are subject-involving: namely, it claims that exemplifying such properties is a matter of bearing the color constitutive relation (whatever that is) to a subject and a circumstance.

That said, the hybrid view envisaged here seems clunky in a few respects that make it not worth accepting. First, if it was supposed to be an advantage of the self-locating property view that it avoids the relationalist's dual levels of properties, the hybrid view gives up this advantage. Second, as noted, the hybrid view contains the apparatus to explain ecumenicism twice over, and this seems redundant. And while this redundancy might be worth accepting if it were the only way of accounting for comparability, we've seen (§2.1, §3.2) that it is not.

²⁵Recall that the latter view delivers ecumenicism not by proliferating properties, but by proliferating centers (viz., $\langle w, t, i \rangle$ triples) with respect to which we can evaluate the extension of each property. But the proliferation of centers is a price everyone should be willing to pay — that's just a consequence of the plurality of worlds, times, and individuals. (Remember that, for present purposes, we can read the commitment to worlds in just as metaphysically inflationary or deflationary terms as we like. Indeed, even if we were forced for some reason to think of a commitment to plural worlds in inflationary terms, presumably the burden of motivating this commitment wouldn't depend on any special features of our account of color, in particular. Hence, the proliferation of world components of centers really is innocuous from the present point of view.)

I don't think so: I suggest that the contemplated ontological explosion is unobjectionable, and provides no grounds for rejecting color relationalism. For even if you thought it objectionably unparsimonious to proliferate non-relational properties (I take no official stand here), it should be much less worrisome to recognize that there are many different relations, hence relational properties, that obtain between individuals. To see the point, note that I bear the relational property of being less than n years old for infinitely many natural numbers $n > 100$; therefore we can fix one of the places in the binary *has lived fewer years than* relation between organisms and numbers to obtain an infinitude of relational properties I bear to numbers. I take it that recognizing these infinitely many relational properties is no serious burden on a theory of age that has already recognized organisms and numbers. But it is only this second, untroubling, kind of proliferation to which relationalists are committed. Thus, the relationalist supposes that we have already recognized the existence of pluralities of objects, subjects, and viewing conditions. Against this backdrop of commitments, the acceptance of a further plurality of relations holding among the objects, subjects, and viewing conditions we have already recognized seems unproblematic (cf. Cohen 2009, pp. 133–136).

A further parsimony-based objection might involve relationalism's dual sets of properties — it recognizes both fine-grained color properties (which it takes to be perceptually represented and which it uses to secure the benefit of ecumenicism) and coarse-grained color properties (which it takes to be cognitively/perceptually represented and which it uses to secure the benefit of comparability). Once again, though this dual commitment may initially seem undesirable, I don't believe this is a crucial worry for the view. Recall that, for the relationalist, both fine-grained and coarse-grained colors are relational properties, constituted in terms of relations to perceivers and perceptual conditions. Though they differ in the specificity or grain of the relation, they are carving up the same underlying reality. Moreover, relationalists argue that both sorts of properties (/ways of cutting the underlying reality) are warranted by their explanatory benefits. As such, I do not regard it as an ultimately telling ontological shortcoming of the view that it works with two levels of relational properties rather than one (cf. Cohen 2009, pp. 114–116).

6.2 Perceptual comparability

Several commentators have objected that the color relationalist story on comparability inappropriately locates all agreement and disagreement at the level of cognitive/linguistic representation, and in particular that it wrongly precludes *perceptual* agreement and disagreement.

Of course, as we have seen, relationalism allows for comparability in the representation of coarse-grained color. But this, at least for some commentators, has seemed insufficient. Thus, for example, Egan writes:

One might have thought (and I would like to be able to think) that when Ernie and Vert are spectrum inverted relative to each other, or when Ron and Hermione differ about whether a given object is unique blue or instead is a bit greenish, there is a representational incompatibility, not just in language — not just in their verbalized color-attributions — but also in perception (Egan 2012, p. 309).

Brogaard presses a version of this complaint as well:

If color experience is thought to represent color properties, however, then the relational view runs into trouble.... It is fairly plausible to think that you and I could have indistinguishable color experiences in the same viewing conditions. Yet if phenomenology determines the representational content of perception, which is also plausible, then it follows that you and I cannot be constituents of the representational content. Yet on Cohen's view, if color properties are constituents of the representational content of perception, then perceivers are also constituents of the representational content of perception. But now we have a contradiction. So, one of the premises must go (Brogaard 2014, p. 12; cf. Brogaard 2012, pp. 317–318).²⁶

And both Egan and Brogaard regard relationalism's inability to provide for comparability at the level of the perceptual representation of color as a reason to prefer a self-locating property theory.²⁷

I agree with these critics that we should want to preserve the possibility of interpersonal (and for that matter, intrapersonal) agreement and disagreement, hence comparability, in the representation of color. I find it much less obvious that there must be agreement and disagreement, hence comparability, in the *perceptual* representation of color. After all, our evidence of interpersonal/intrapersonal agreement and disagreement (e.g., intrapersonal matching behavior, interpersonal verbal ascriptions) appears to draw on representations occurring at both perceptual and non-perceptual levels (e.g., judgment and linguistic levels of representation). Given this fact, it seems inappropriate to rule out explanations of comparability framed in terms of the non-perceptual levels of representation we have independent reason to believe are at work, such as that offered by relationalists.²⁸

6.3 Hard comparability cases

A third cluster of concerns for relationalism involve hard cases of comparability that might seem to resist treatment in terms of the relationalist apparatus provided so far.²⁹

6.3.1 Intercontextual comparability

Recall that the relationalist account of comparability involved supplementing the overtly expressed content of color ascriptions with values for perceiver and perceptual circumstance parameters that are fixed by the contexts in which those ascriptions are

²⁶Indeed, Brogaard (2014, p. 12) points out that if we insist on the possibility that color experiences represent interpersonally shareable contents, relationalists will be forced to conclude that color experiences cannot represent color properties (even though they can represent shape, texture, and other visually accessible properties). I agree with Brogaard that this conclusion is implausible, and see it as yet another reason that relationalists should deny Brogaard's assumption that the contents of perception/perceptual experience must be shareable.

²⁷Other critics who have advanced versions of the same objection, though not in the service of motivating the self-locating property theory, include Allen (2012, p. 16, 2011, p. 318); Pautz (2010); Tye (2012, p. 12).

²⁸Analogy: It's fair to demand that an adequate overall account of language understanding should predict the unacceptability of (6).

(6) a. Buffalo buffalo Buffalo buffalo buffalo buffalo Buffalo buffalo. (Pinker 2000, p. 210)

But it's not fair to demand that an adequate syntactic theory, in particular, should be the particular component of the overall account of language understanding from which the prediction is derived.

²⁹Thanks to Adam Pautz for pushing me on these issues.

uttered or tokened. This treatment allows us to understand how the hidden parameters in ascriptions made by two speakers within the very same context could be assigned the very same (coarse-grained) values, so that the resulting contents could be comparable after all. At first blush, however, it is not obvious how to extend this treatment to cases of *intercontextual* agreement/disagreement, where speakers in distinct contexts ascribe colors to one and the same object.

Consider, for example, the following sequence of ascriptions, where the first is made by speaker S_1 in perceptual circumstance C_1 and conversational context K_1 , and the second is made by speaker S_2 perceptual circumstance C_2 and conversational context K_2 :

- (7) a. S_1 in C_1 in context K_1 : a is unique green,
 b. S_2 in C_2 in context K_2 : a is not unique green/is bluish green.

As we have seen, relationalism predicts that these ascriptions ascribe distinct, and incomparable, fine-grained properties. However, if we move to the coarse-grained contents the relationalist uses to account for comparability, we find that the properties attributed by these two ascriptions are, respectively:

- (8) a. *unique green to a K_1 -relevant perceiver in K_1 -relevant perceptual circumstances,*
 b. *not unique green to a K_2 -relevant perceiver in K_2 -relevant perceptual circumstances (bluish green to a K_2 -relevant perceiver in K_2 -relevant perceptual circumstances).*

Of course, there is no reason to expect, in general, that what is K_1 -relevant will overlap with what is K_2 -relevant, and so reason to expect, in general, that K_1 and K_2 will agree about what counts as the contextually relevant type of perceiver and circumstance (cf. note 15). But if not, then it would seem relationalists will be unable to account for intuitions to the effect that ascriptions like (7a–b), occurring in different conversational contexts, are ever comparable (hence can agree or disagree).

But this is too quick. First, though I said that *in general* we can't expect overlap between what is K_1 -relevant and what is K_2 -relevant, there certainly can be cases where there is overlap of that kind, and where, consequently, understanding (7a–b) as expressing (8a–b) allows for intercontextual comparability. Indeed, presenting (7a–b) in such close proximity, as I have, plausibly encourages us to think we're in such a situation, and therefore may account for such intuitions of comparability as there are. Second, note that when it is not assumed that the contexts overlap in what perceivers and perceptual conditions they make relevant, the intuitions about comparability are weak: it's not all that clear that (7a–b) really are comparable in the intended sense when set in such different conversational contexts. And third, while the relationalist apparatus predicts that ascriptions like (7a–b), when made in distinct contexts that disagree about the relevant perceivers/perceptual circumstances, do not agree in the content they express, there is a weaker sense of agreement that the view can accommodate even with respect to such pairs of contexts. Namely, (7a–b) are such that they *would have* expressed comparable (indeed, conflicting) contents had they, counterfactually, been interpreted relative to a common conversational context.³⁰ If you

³⁰Analogy: Consider the following pair of ascriptions:

- (9) a. S_1 : Jones is tall.
 b. S_2 : Jones is not tall.

Presenting these ascriptions in such close proximity encourages us to treat the two ascriptions as occurring in a single context, and to take them as comparable/conflicting. When we add that

like, this is a notion of comparability/conflict that applies at the level of the character rather than content. My present point is that, to the extent there remains a weak intuition of comparability with respect to contexts that disagree about what perceivers and perceptual circumstances are relevant, that intuition might reflect comparability at the level of character rather than content. If so, then such cases present no obstacles for color relationalism.³¹

Parallel issues arise in the consideration of intercontextual disquotational reports of color ascriptions, such as (11a–b):

- (11) a. S_1 in C_1 in context K_1 : a is unique green.
 b. S_2 in C_2 in context K_2 : S_1 says that a is unique green.

Here the worry is that if, as before, K_1 -relevant perceivers/circumstances are different from K_2 -relevant perceivers/circumstances, then the content expressed in (11a) won't match/be comparable with that expressed by the embedded clause in (11b). And this would appear to threaten the possibility of (true) intercontextual disquotational reports of color ascriptions. (The problem might look even worse in the special case where $S_1 = S_2$.)

Once again, I suggest that these difficulties are exaggerated. First, it is useful to remind ourselves that many (most?) ordinary instances of color discourse between normal trichromatic human beings occur in contexts that (though perhaps differing in other semantically important ways) *do* overlap in what sorts of perceivers/circumstances they make relevant to our color ascriptive/explanatory needs. (There are, of course, many fine-grained differences that won't emerge in ordinary conversational settings.) And disquotation between pairs of contexts that overlap in that way present no difficulties for the view. Of course, there are also conversational contexts in which we engage in (formal or informal) contemplation of perceptual variation with respect to color (of the

S_1 is a high school basketball coach and S_2 is an NBA scout, we add information that makes it at least less likely (but not impossible; see note 33) that the two utterances occur in contexts that overlap in the standard of comparison for 'tall' that they make available; consequently, the impression of comparability weakens significantly. Interestingly, however, even when the impression of comparability weakens in this way, there remains a modest sense of conflict between (9a–b), plausibly arising from our knowledge that the two linguistic forms can be used to generate mutually impossible contents by evaluating them with respect to a common context, even if they were not so evaluated on this occasion.

³¹Objection: The weak, character-based sense of comparability/conflict is *too* weak. After all, the same, weak, character-based sense of comparability applies to John's and Mary's utterances in (10a–b):

- (10) a. John: I am a doctor.
 b. Mary: I am not a doctor.

But surely there is no impression that John's and Mary's utterances disagree — certainly not in the sense that the disagreement data of §1.2 present, and that we took it as (part of) our brief to account for.

Response: The character-based sense of comparability/conflict at issue is indeed weak (though, I believe, non-vanishing), and would not be plausible if offered as a full account of comparability intuitions with respect to color attributions, generally speaking. But the relationalist does not offer it as a full account of comparability intuitions with respect to color attributions, generally speaking. On the contrary, she claims that there are plenty of color attributions that are comparable in a much more robust sense (*viz.*, those made in contexts that agree about the perceivers/perceptual circumstances they make relevant). The present point is that even where relationalism does not allow for conflict in that robust sense, there is a weaker sense of conflict that remains, and that may be the source of any residual intuitions of comparability that might otherwise be adduced as evidence that the account undergenerates cases.

sort, for example, that we saw in §1.1). But it is much less obvious that intercontextual disquotation involving those sort of contexts goes through straightforwardly, so correspondingly doubtful that we should want a theory predicting that it does. Finally, I would point out that, as before, relationalists can account for whatever intuitions of comparability remain in terms of the weak, character-based level of comparability: we can treat the embedded clause of S_2 's report in (11a) as agreeing with/comparable to S_1 's ascription in (11b) in the contents they *would have* expressed if, counterfactually, they had been interpreted relative to a common context. For these reasons, I don't take the objections based on intercontextual disquotational ascription reports to be ultimately damning for color relationalism.

6.3.2 Intracontextual comparability

There is one final objection based on comparability that has to do with intracontextual differences arising in perceivers that are (though different) both plausibly K -relevant, and who therefore pull our understanding of the relationalist's level of comparability-ensuring contents in incompatible directions. Thus, consider two subjects S_1 and S_2 , who both pass standard (if stipulative) psychophysical tests for normal color vision, and who are in (possibly different) perceptual circumstances C_1 and C_2 , which both pass ordinary (if stipulative) psychophysical standards for normality. And now suppose these two subjects enter into a single conversational context K , in which they proceed to discuss the color of an individual a by making the following sincere, perceptually informed, ascriptions:

- (12) a. (Normal) S_1 in (normal) C_1 in context K : a is unique green.
 b. (Normal) S_2 in (normal) C_2 in context K : a is not unique green/is bluish green.

(This can certainly happen, given the facts about interpersonal variation in color perception discussed in §1.1.) We've said that relationalists will attempt to account for the apparent conflict/comparability of these (12a–b) by treating them as expressing the conflicting/comparable contents (13):

- (13) a. *unique green to a K -relevant perceiver in K -relevant perceptual circumstances,*
 b. *not unique green to a K -relevant perceiver in K -relevant perceptual circumstances*
(bluish green to a K -relevant perceiver in K -relevant perceptual circumstances).

How should we understand that pair of contents? And, in particular, given the differences between our two different (stipulatively normal) interlocutors and their (stipulatively normal) perceptual circumstances in the common conversational context K , what sort of perceiver/circumstance should we say is K -relevant? Of course, the whole point of invoking coarse-grained colors in the first place was to allow us to count *both* S_1/C_1 and S_2/C_2 as falling under the parametric specifications that contextualists claim ordinary color ascriptions tacitly harbor. That both S_1/C_1 and S_2/C_2 are (by stipulation) normal only adds to the plausibility of this choice. But if that's so, and if S_1 and S_2 speak truly in uttering (13a–b), then either a has incompatible properties, or we lose the conflict/comparability desideratum relationalists were trying to secure with this level of content: we would be saying that a exemplifies unique green relative to one type of K -relevant perceiver/circumstance, and exemplifies not unique green/bluish green to a different type of K -relevant perceiver/circumstance.³²

³²I should note that there are at least some cases with this structure in which the comparability intuition can become extremely weak; of course, relationalism is not challenged by those cases.

This is an instance of a much more general worry that arises for any form of contextualism — any view on which the interpretation of linguistic material is sensitive to a contextually supplied parameter. The worry arises when we try to evaluate multiple occurrences of such contextually-sensitive linguistic forms within a single context, as, for example, when multiple (and equally admissible) speakers in a single context apply and forbear a gradeable adjective such as ‘tall’ to a single individual. In terms of a well-known metaphor from Lewis (1979b), the worry here is that contextualists will be forced to interpret context-sensitive expressions within a discourse by appeal to multiple, non-equivalent scoreboards rather than a single, discourse-wide scoreboard, and so will be forced to conclude that the distinct tokens of the linguistic forms in question will be *differently* supplemented by context (i.e., according to the values on different scoreboards), hence will express distinct/incomparable contents (cf., e.g., Richard 2004, pp. 215–216).

I think the predicament just outlined is predictable, and not fatal to contextualism (either generally speaking or for the specific form of contextualism I have appealed to in defending color relationalism). It is unsurprising that when there is just a single admissible candidate value for the relevant contextual parameter (e.g., a single speaker’s assumed threshold for tallness, a single speaker’s assumptions about what kinds of perceivers/circumstances are normal), hearers feel pressure to accommodate by interpreting the context-sensitive expressions in a way that adverts to that single admissible candidate value — hence we can obviously model the discourse with a single scoreboard. It is also unsurprising that when there are multiple admissible candidate values for the relevant contextual parameter (e.g., assumed thresholds for tallness of the different ascribers of ‘tall’, the assumptions about what kinds of perceivers/circumstance are normal made by every speaker who uses a color predicate in the shared context), hearers feel torn between the accommodative/interpretive alternatives. Therefore, it should be similarly unsurprising, given contextualism, that these will be hard cases — viz., cases in which it is not obvious which, if any, single scoreboard is semantically relevant.

But, as DeRose (2004) argues persuasively, it just does not follow from this unobviousness that contextualists are forced to accept a multiple scoreboard model: it does not follow that contextualists are committed to evaluating any pair of distinct occurrences of context-sensitive expressions as being separately/differently supplemented by context. On the contrary, it is fully compatible with the commitments of contextualism that there should be a procedure for installing a single scoreboard within a discourse, such that different tokens of a given context-sensitive expression-type within that discourse receive the very same contextual supplementation — even when the multiplicity of admissible candidate values available in the discourse place *prima facie* conflicting demands on interpreters.³³

What follows is an attempt to explain how relationalists/contextualists can secure comparability in cases where it is needed. If it should turn out that there are few such cases at the end of the day, then no harm done.

³³ Just what sort of a single-scoreboard semantics should we accept? There are many options here, and different contextualists will have different preferred answers, and possibly different answers for different kinds of cases. DeRose (2004), who is defending contextualism about knowledge attributions, rather than color attributions, considers a (non-exhaustive) range of single-scoreboard rules for contextual setting of epistemic standards including the following: higher standards prevail, non-vetoed standards prevail, a binding arbitration model, community deference, an ‘exploding scoreboard’ rule (extensive semantic gappiness for knowledge attributions), more limited gappiness rules of various kinds, and a supervaluationist

I conclude, therefore, that relationalists can indeed make sense of intracontextual comparability, even in conversational contexts where different perceptual reports make ostensibly competing demands on our understanding of which perceivers/circumstances are relevant.

7 Conclusion

This paper has (mostly) been devoted to a comparative assessment of two competing color ontologies designed to respect both of the desiderata of §1. As noted earlier, it is possible to resist the broadly empirical cases for either or both of those desiderata. But for those who find the motivations for the desiderata compelling, it will be seen as a virtue of both color relationalism and the self-locating property view that they are poised to deliver both.

The two views under consideration share enough — and enough that distinguishes them from alternative accounts on offer — that it is reasonable to regard them as near neighbors, and the choice between them as somewhat internecine. Nonetheless, I believe there are significant differences between the two accounts, and that they are consequently vulnerable to different kinds of concerns. In particular, it seems to me that the self-locating property view faces worries about its scope, motivation, and the limits of its systematic explanatory power, whereas the most important questions relationalism faces are smaller bore (but still important) issues about the extravagance of its ontological commitments, its non-perceptual account of comparability, and the ability of its contextualist semantics to account for various hard cases.

This difference in the problems faced by the two views means that the considerations above don't amount to a clear apples to apples comparison between the accounts. And, of course, different theorists will have different views about the relative significance of the two classes of problems. However, I find a further asymmetry in the problems faced by the two accounts: while I don't know how the problems for the self-locating view can be answered satisfactorily, I have argued, above, that the most important challenges to color relationalism can be met.

I do not claim to have settled once and for all the dispute between color relationalism and the self-locating property view of color. But I hope I have contributed to a clearer understanding of them and their relative strengths and weaknesses.³⁴

rule. I see no reason versions of these proposals couldn't be adapted to the sort of contextualism about ordinary color attributions we are considering here.

(Perhaps needless to say, if there is such a procedure for installing a single scoreboard for contextual supplementation within a discourse, there's no reason to think of that procedure as marking metaphysically significant distinctions; hence, the idea that there is a single conversational scoreboard in no way supports the view that there is a metaphysically significant asymmetry between perceptual variants in the context of the argument from perceptual variation discussed in §1.1.)

³⁴Thanks to Berit Brogaard, Daniel Burnston, Mazviita Chirimuuta, Damon Crockett, Andy Egan, Matthew Fulkerson, and an anonymous referee for this journal for comments and discussion, and to audiences at Auburn University, the University of Glasgow, and the University of London, who heard earlier versions of this material.

References

- Allen, Keith (2011), "The Red and the Real: An Essay on Color Ontology, By Jonathan Cohen," *European Journal of Philosophy*, vol. 19, 2, pp. 315-318.
- (2012), "Colour, Contextualism, and Self-Locating Contents," *Croatian Journal of Philosophy*, vol. 12, 3, pp. 331-350.
- Armstrong, David M. (1987), "Smart and the Secondary Qualities," in *Metaphysics and Morality: Essays in Honour of J. J. C. Smart*, ed. by Philip Pettit, R. Sylvan, and J. Norman, Reprinted in Byrne and Hilbert 1997, 33-46, Basil Blackwell, Oxford.
- Boghossian, Paul A. and J. David Velleman (1989), "Colour as a Secondary Quality," *Mind*, vol. 98, Reprinted in Byrne and Hilbert 1997, 81-103., pp. 81-103.
- Brogaard, Berit (2012), "Color Eliminativism or Color Relativism?" *Philosophical Papers*, vol. 41, 2, pp. 305-321.
- (2014), "The Self-Locating Property Theory of Color," *Minds and Machines*.
- Byrne, Alex (2009), "Experience and Content," *Philosophical Quarterly*, vol. 59, 236, pp. 429-451.
- Byrne, Alex and David R. Hilbert (eds.) (1997), *Readings on Color, Volume 1: The Philosophy of Color*, MIT Press, Cambridge, Massachusetts.
- (2003), "Color Realism and Color Science," *Behavioral and Brain Sciences*, vol. 26, 1, pp. 3-64.
- Caponigro, Ivano and Jonathan Cohen (2011), "On Collection and Covert Variables," *Analysis*, vol. 71, 3, pp. 478-488.
- Chalmers, David (2006), "Perception and the Fall From Eden," in *Perceptual Experience*, ed. by Tamar Szabó Gendler and John Hawthorne, Oxford University Press, New York, pp. 49-125.
- Cohen, Jonathan (2009), *The Red and The Real: An Essay on Color Ontology*, Oxford University Press, Oxford.
- (2012), "Redder and Realer: Responses to Egan and Tye," *Analytic Philosophy*, vol. 53, 3, pp. 313-326.
- Dancy, Jonathan (1986), "Two Conceptions of Moral Realism," *Proceedings of the Aristotelian Society*, vol. Supplementary Volume 60, pp. 167-187.
- DeRose, Keith (2004), "Single Scoreboard Semantics," *Philosophical Studies*, vol. 119, 1/2, pp. 1-21.
- Egan, Andy (2006), "Secondary Qualities and Self-Location," *Philosophy and Phenomenological Research*, vol. 72, 1, pp. 97-119.
- (2010), "Disputing About Taste," in *Disagreement*, ed. by Ted Warfield and Richard Feldman, OUP.
- (2012), "Comments on Jonathan Cohen's *The Red and the Real*," *Analytic Philosophy*, vol. 53, 3, pp. 306-312.
- Gibbard, Allan (2006), "Moral Feelings and Moral Concepts," in *Oxford Studies in Metaethics, 1*, ed. by Russ Schafer-Landau, Clarendon Press, Oxford, pp. 195-215.

- Hazlett, Allan and Edward Wilson Averill (2010), "A Problem For Relational Theories of Color," *Philosophy and Phenomenological Research*, vol. 81, 1, pp. 140-145.
- Hurvich, Leo Maurice (1981), *Color Vision*, Sinauer Associates, Sunderland, Massachusetts.
- Jacobs, Gerald H. (1981), *Comparative color vision*, Academic Press series in cognition and perception, Academic Press, New York.
- Jitsumori, Masako and Juan D. Delius (2001), "Object Recognition and Object Categorization in Animals," in *Primate Origins of Human Cognition and Behavior*, ed. by Tetsuro Matsuzawa, Springer Japan, Tokyo, pp. 269-293.
- Kaplan, David (1989), "Demonstratives: An Essay on the Semantics, Logic, Metaphysics, and Epistemology of Demonstratives and Other Indexicals," in *Themes From Kaplan*, ed. by Joseph Almog, John Perry, and Howard Wettstein, Oxford University Press, New York, chap. 17, pp. 481-563.
- Kuehni, Rolf G. (2004), "Variability in Unique Hue Selection: A Surprising Phenomenon," *Color Research and Application*, vol. 29, pp. 158-162.
- Levin, Janet (2000), "Dispositional Theories of Color and the Claims of Common Sense," *Philosophical Studies*, vol. 100, pp. 151-174.
- Lewis, David (1979a), "Attitudes *De Dicto* and *De Se*," *The Philosophical Review*, vol. 88, pp. 513-543.
- (1979b), "Scorekeeping in a Language Game," *Journal of Philosophical Logic*, vol. 8, pp. 339-59.
- Matthen, Mohan (1999), "The Disunity of Color," *The Philosophical Review*, vol. 108, 1, pp. 47-84.
- (2005), *Seeing, Doing, and Knowing: A Philosophical Theory of Sense Perception*, Oxford University Press, Oxford.
- Pautz, Adam (2010), "An Argument for the Intentional View of Visual Experience," in *Perceiving the World*, ed. by Bence Nanay, Oxford University Press, New York.
- Peacocke, Christopher (1979), "Deviant Causal Chains," *Midwest Studies in Philosophy*, vol. 4, pp. 123-155.
- Perry, John (1979), "The Problem of the Essential Indexical," *Noûs*, vol. 13, pp. 3-21.
- Pinker, Steven (2000), *The Language Instinct: How the Mind Creates Language*, Harper Perennial Modern Classics, New York.
- Richard, Mark (2004), "Contextualism and Relativism," *Philosophical Studies*, vol. 119, 1-2, pp. 215-242.
- Siegel, Susanna (2010a), "Do Experiences Have Contents," in *Perceiving the World*, ed. by Bence Nanay, Oxford University Press, New York.
- (2010b), *The Contents of Visual Experience*, Oxford University Press, New York.
- Smith, A. D. (2002), *The Problem of Perception*, Harvard University Press, Cambridge, Massachusetts.
- Soto, Fabian A. and Edward A. Wasserman (2010), "Comparative Vision Science: Seeing Eye to Eye?" *Comparative Cognitive Behavioral Review*, vol. 5, pp. 148-154.

- Spetch, Marcia L. and Alinda Friedman (2006), "Comparative cognition of object recognition," *Comparative Cognition & Behavior Reviews*, vol. 1, pp. 12-35.
- Stalnaker, Robert (1979), "Assertion," in *Syntax and Semantics 9: Pragmatics*, ed. by Peter Cole, Academic Press, New York, pp. 315-332.
- Thompson, Evan (1995), *Colour Vision: A Study in Cognitive Science and the Philosophy of Perception*, Routledge, New York.
- Thompson, Evan, Adrian Palacios, and Francisco Varela (1992), "Ways of Coloring: Comparative Color Vision as a Case Study for Cognitive Science," *Behavioral and Brain Sciences*, vol. 15, pp. 1-74.
- Travis, Charles (2004), "The Silence of the Senses," *Mind*, vol. 113, 449, pp. 57-94.
- Tye, Michael (2012), "Cohen on Color Relationalism," *Analytic Philosophy*, vol. 53, 3, pp. 297-305.