

# Redder and Realer: Responses to Egan and Tye

Jonathan Cohen\*

*We lament the mistakes of a good man, and do not begin to detest him until he affects to renounce his principles.*

— Letters of Junius, No. 41, to Lord Mansfield (14 November 1770)

I am grateful to Egan and Tye for their thoughtful and challenging criticisms of *The Red and the Real* (henceforth *RR*), from which I have learned much. They raise a number of interesting issues, and I am thankful for the opportunity their criticism provides to reconsider some of the key commitments of the book and attempt to clarify my thinking about these matters. While reasons of space prevent me from taking on all of their criticisms, in what follows I'll respond to what I take to be the most important issues they raise.

## From Variation to Relationalism?

I begin with Tye's criticisms of the argument from perceptual variation, which is the most important motivation for color relationalism offered in *RR*.<sup>1</sup> Roughly this (non-deductive) argument begins with the empirical claim that there is significant inter- and intra-personal variation in representational responses to a given color stimulus, and the thought that, on standard assumptions, each variant represents the color of the stimulus. The next step is a symmetry claim — viz., that in such cases, it is hard to imagine what, other than *ad hoc* stipulation, could make it the case that just one of the variants is uniquely veridical (i.e., veridical at the expense of the other variants). The argument then appeals to the principle that we should avoid *ad hoc* stipulation

---

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

<sup>1</sup>Since both Tye (2012) and Cohen (2012) present this argument in some detail, I confine myself here to an especially compressed characterization. My fullest presentation of the argument, of course, occurs in *RR* itself.

where possible, and takes this as reason for preferring an ecumenical view on which the ostensibly competing variant representations of the stimulus's color can be simultaneously veridical. Finally, the argument involves an abductive inference to the relationalist view that colors are constituted in terms of relations to perceivers and viewing conditions on the ground that such a color relationalism gives us the best way of implementing the desired ecumenicism.

Tye is unmoved by this argument. In the face of the observed inter- and intra-personal variation in representational responses to a given color stimulus, Tye rejects the ecumenical approach outlined above, and instead favors a kind of color epistemicism on which there must be one variant that veridically represents the color of the stimulus (at the expense of the others) even if we don't know which it is (Tye (2012, 2–3); cf. Byrne and Hilbert (2003, 17), Byrne and Tye (2006)).

I agree with Tye that this epistemicism is a coherent response to cases of representational variation. In fact, it seems to me there are some cases where it is not only coherent, but obviously the right view of the situation. For example, if you and I disagree about whether the liquid in a particular glass is water, or if we disagree about whether some particular plane figure is a square, the epistemicist line seems clearly correct: there is a fact of the matter about which of our representations is veridical at the expense of the other, whether or not anyone knows that fact (now or ever). On the other hand, there are also cases of representational variation about which the kind of epistemicism at issue is pretty obviously the wrong view of the situation. Thus, if you and I disagree about whether some particular joke is funny, or if we disagree (while occupying different reference frames) about whether a particular object is moving at 60mph, it seems absurd to think that there is a fact of the matter about which of our representations is veridical at the expense of the other, whether or not anyone knows this fact (now or ever).

It appears, then, that there are both paradigm instances of representational variation in which epistemicism seems obviously preferable to ecumenicism and paradigm instances of representational variation in which ecumenicism seems obviously preferable to epistemicism. As such, the mere existence of these classes of cases leaves unresolved the central question of how we should treat the disputed case of representational variation with respect to color. That is, one can't resolve that central question by pointing to either class without offering some reason for thinking that the variation about color should be assimilated to that class rather than the other. And this is just to say that what Tye calls the "odor of verificationism" attaching to ecumenicist responses to instances of variation doesn't always stink, and that we are crucially in need of a way to choose between it and what we might think of as an "odor of dogmatism" as demanded by individual cases.

Tye is aware of the need for such a method of choosing, and forthrightly gives us his preferred method for making the choice: a priori reflection. He writes that,

A priori reflection tells us that nothing can be both red and green all over at the same time. This is not implicitly relative to a single type of perceiver and a circumstance (I would say), at least as it is usually understood, any more than is the claim that nothing can be both round and square at the same time. This being the case, if grass looks red to [ $S_1$  in  $C_1$ ] and green to [ $S_2$  in  $C_2$ ] then grass cannot be *both* as it appears to [ $S_1$  in  $C_1$  and  $S_2$  in  $C_2$ ]. At most grass is one of the two. Generalizing, a priori reflection provides us with a reason for supposing that, *for each colored thing*, there is a single variation under which the thing has the color it then looks to have (3).

But Tye's appeal to a priori reflection at this juncture is deeply unsatisfying.

First, that method is unreliable in parallel situations. After all, we can easily imagine someone responding to what the special theory of relativity tells us about motion ascriptions by saying that a priori reflection tells us that nothing can be both in motion and at rest at the same time, that this is not implicitly relative to a single reference frame, hence that at most one of a pair of ostensibly competing motion ascriptions must be correct. I take it that this response does not give us reason for rejecting what STR (*qua* physically motivated, broadly empirical theory of the nature of motion properties) tells us about its target, and suggest that the same lesson applies to Tye's exactly analogous remarks in response to color relationalism (*qua* psychophysically motivated and broadly empirical theory of the nature of color properties).

A second point to note about Tye's appeal to a priori reflection is that in making it he explicitly recognizes, but then rejects without argument an alternative analysis of the exclusion phenomenon cited in its first sentence.<sup>2</sup> This reanalysis involves the idea that, while relationalism allows that  $b$  can simultaneously and throughout its extent exemplify both *red for  $S_1$  in  $C_1$*  and *green for  $S_2$  in  $C_2$* , we can nonetheless recapture the exclusion between red and green by noting that nothing can simultaneously exemplify throughout its extent both *red for  $S_1$  in  $C_1$*  and *green for  $S_1$  in  $C_1$* .<sup>3 4</sup>

---

<sup>2</sup>Though Tye doesn't say a lot about his reasons for dissatisfaction with this analysis, he seems to be relying mainly on a comparison between the case at hand and the situation of representational variation about shape (about which such an alternative reanalysis of the exclusion claims would presumably be incorrect). But that case, by itself, can't motivate a view about how to understand representational variation (/exclusion) about color, since (again) it is exactly up for grabs at this point whether that sort of case (as opposed to the sort of case involving funny/motion/etc.) is the right model for thinking about color.

<sup>3</sup>For discussion of this, and other, ways of accommodating such exclusion relations, see RR, 80–81. See also note 6, below.

<sup>4</sup>At one point (3) Tye attempts to buttress the epistemicist reaction to perceptual variation (and perhaps also his favored analysis of the facts of color exclusion) by embracing Pryor's "dogmatic" response to the Cartesian skeptic — roughly, the view that we are warranted in retaining our ordinary belief that we live in a material world (etc.) in the absence of a good reason to abandon that belief, even if we lack a satisfying refutation of the skeptic. While these issues deserve more space than I can give them, I would point out that, to serve its purpose at this point in the dialectic, Tye needs to extend his dogmatism to not only the relatively uncontroversial claims about color rooted in quotidian experience (e.g., claims to the effect that ordinary things exemplify colors, that red and green stand in some kind of exclusion relation, etc.), but also to the correctness of his

As far as I can see, then, the considerations Tye adduces fail to make an advance on the problem of choosing whether to assimilate instances of representational variation with respect to color to the model of variation about shape or natural kinds (where an epistemicist response seems preferable) or to the model of variation about humor or motion (where an ecumenical response seems preferable). However, I do not believe that the situation is hopeless; it is for this reason that I offer a number of considerations in *RR* (esp. 46–53) that, though non-demonstrative, are intended to move past the impasse.

Perhaps most importantly, it should be emphasized that the inference from variation to ecumenicism, as presented in *RR*, does not rest solely on the bare existence of perceptual variation. On the contrary, that inference depends on both the facts about variation and, crucially, on an additional inductive (viz., non-deductive) inference from the failure of the best-motivated attempts to single out a uniquely veridical variant in particular cases of perceptual variation. In particular, in laying out the argument, I consider various kinds of perceptual variation with respect to color in some detail, and argue that, in case after case, the most promising attempts to single out a uniquely veridical variant — attempts resting on industrial and scientific standards reflecting genuine epistemic needs — require unmotivated stipulations. Significantly, the considerations underlying these verdicts about the relevant industrial and scientific standards are local to the particulars about color perception. This matters because, given the availability of both epistemicist and ecumenicist models for thinking about representational variation, we should be wary of extending morals from very different kinds of cases to the one in dispute.<sup>5</sup>

I am emphasizing the non-demonstrative character of these considerations (as I did in *RR*) because that is to be expected in such broadly empirical settings; like it or not, this is how investigation proceeds outside mathematics. It is, as

---

preferred philosophical analyses of those uncontroversial claims (e.g., that the exclusion relation between red and green is not to be understood as relative to a type of perceiver and circumstance). While I can at least imagine why someone might endorse dogmatism about at least much of what ordinary experience tells us about colors, I have a hard time seeing how ordinary experience is committed one way or the other about such matters of philosophical analysis. Hence, I don't see that Pryor's sort of dogmatism will do for Tye what he wants it to do.

<sup>5</sup>By contrast, Tye's comparison to the bank caper case, wherein "I do not have excellent reason to believe *of* any particular one of [A, B, and C] that *he* murdered the bank teller but I do have excellent reason for believing *that* one of them murdered the bank teller" seems disanalogous in key respects. Most significantly, in Tye's case our independently extremely well-confirmed views about how objects move, human psychology, and the like (plus the setup of the case) suggest that it is overwhelmingly likely that there is a single murderer, while our general beliefs about the world make it thoroughly understandable how there could be a murderer without our having sufficient evidence of her guilt. The first of these beliefs creates the presumption that there must be a culprit among the three, and the second explains away potential counter-considerations to the presumption resulting from our lack of evidence about specific individuals. But this situation contrasts markedly with the case of perceptual variation about color, where our background beliefs about object movement and the like are insufficient either to install a presumptive prejudice in favor of a uniquely veridical variant or to explain away counter-considerations resulting from our lack of evidence. If anything, we have substantial (but defeasible) reason to believe that there is not uniquely veridical variant: here the failure of several hundred years of systematic efforts directed at uncovering such facts of the matter arguably establishes a presumptive case against their existence.

Tye says, metaphysically possible that there is an unknown fact of the matter about which variant is veridical at the expense of the others. But while this *could* be true about perceptual variation, there is no reason to suppose that it is, and examination of the best motivated and most clearly relevant evidence suggests that it is not. Possibility does not probability or plausibility make.

## Relationalism and Disagreement

Tye raises the further worry that the relationalist's commitment to ecumenism about perceptual variation leaves her unable to account for intrapersonal and interpersonal disagreement about color. Of course, the relationalist's ecumenical response to perceptual variation is designed to render  $S_1$ 's and  $S_2$ 's color attributions compatible — even when their verbal expression suggests otherwise — by construing the content of those attributions as involving distinct, compatible properties. But if we secure compatibility by holding that  $S_1$ 's and  $S_2$ 's attributions don't attribute and forebear the very same property, then Tye objects that it's hard to see how  $S_1$  and  $S_2$  could ever make incompatible color attributions, hence how any two color attributions could be in disagreement.

My strategy of response to this concern in *RR* involves holding that cognitive and verbal representations of color attribute properties that are coarse-grained than the fine-grained properties attributed by the visual system. Thus, for example, I claim that a thought/utterance of 'this ripe lemon is yellow' in context  $K$  attributes to a contextually indicated lemon the "coarse-grained" color *yellow for the perceivers relevant in context  $K$  under the perceptual circumstances relevant in  $K$*  (*RR*, 100). The thought is that even if the distinct fine-grained colors *unique yellow for  $S_1$  in  $C_1$* , *greenish-yellow for  $S_2$  in  $C_2$*  attributed by  $S_1$ 's and  $S_2$ 's visual systems are compatible, we can explain the color disagreement between  $S_1$  and  $S_2$  by claiming that  $S_1$  verbally attributes, while  $S_1$  verbally forebears, a common coarse-grained color (viz., one whose context parameter is assigned to the very same value).

So supplemented, relationalism has the capacity to describe individuals as disagreeing about the colors of objects (and also to describe the representations underlying those disagreements as involving the very same property) — it is just that the descriptions in question will involve representations of coarse-grained properties at the cognitive/linguistic level rather than representations of the fine-grained properties by the visual system.<sup>6</sup> Therefore, the objection

---

<sup>6</sup> Supplementing relationalism with this contextualist semantics also (among other advantages discussed in *RR*, 122–132) provides yet another way of accounting for the exclusion intuition — the intuition that something's being red is in some sense incompatible with, or excludes, that thing's being green — that, as we saw in §, Tye attempts to press against relationalism. For that intuition is naturally expressed by an utterance, in context  $K$ , of (something like) the string "nothing can be red and green all over at the same time". Now, according to the contextualist semantics just offered, this utterance is true just in case nothing is (all over and at the same time) both red for the  $K$ -relevant perceivers under the  $K$ -relevant perceptual circumstances and green for the  $K$ -relevant perceivers under the  $K$ -relevant perceptual circumstances. But the relationalist will hold

that relationalism doesn't allow for interpersonal and intrapersonal disagreement about color cannot be correct.<sup>7</sup> If there is an objection in the vicinity, then, it must be that the proposed treatment is somehow inadequate. For example, I can imagine someone thinking that the coarse-grained level where the relationalist locates disagreement is not the right explanatory locus. But if that is the objection, then it is not an objection about the bare fact of disagreement, but about what part of the theoretical apparatus is best suited to a particular bit of explanatory labor. And in any case, before we can take this sort of objection seriously, the anti-relationalist owes us an argument against carrying out the explanatory labor in the way that I have proposed.

## Color Illusion, Color Constancy

Tye objects that color relationalists are unable to provide an adequate account of either intrapersonal color illusion or color constancy; since these complaints are intimately related, I will treat them together.

Consider illusion first. Tye points out that the standard description of Adelson's checkershadow illusion (fig. 1) has it that the two different regions labeled *A* and *B* erroneously look to be different in color, when in fact they exemplify the very same color. He correctly points out that relationalists will reject this standard description, and instead will say that *A* and *B* do

---

that this is in fact correct: nothing does exemplify all over and at the same time both of those relational coarse-grained properties (even though she holds that there are things that are, all over and at the same time, both red for  $S_1$  in  $C_1$  and green for  $S_2$  in  $C_2$ ).

<sup>7</sup>Tye seems to miss this point in claiming that my view is unable to handle the case of interpersonal disagreement he presents on p4, where the source of the disagreement is an "unusual physiological condition" affecting one but not the other of the two observers; in fact I describe and explain my treatment of this sort of case in *RR*, 127–128.

There are more problematic cases of interpersonal disagreement, raised by Pautz (2010), in which there are two different subjects, both normal by standard psychophysical criteria, one of whom utters 'chip *a* is unique green' and one of whom utters 'chip *a* is bluish green'. This case is more subtle because it's much less clear that there is a single contextually relevant perceiver type (given the difference between the two psychophysically normal speakers) with respect to which we can assign coarse-grained color contents to the two utterances so as to make them come out as conflicting.

While there's much to say about this more complicated sort of case and the lessons it offers, I'm inclined to take it not as an indictment of contextualism about color contents, but as showing that the ways in which contextual standards are set can be difficult to understand in hard cases. When there's only a single speaker making color attributions who is within whatever standards of psychophysical normality are in play, then there is pressure on evaluators to accommodate as far as we can – i.e., to find admissible values for relevant contextual parameters with respect to which we can evaluate her utterances as being true. By contrast, when there are distinct, psychophysically distinguishable speakers, both normal by standard tests, making such apparently conflicting color attributions in a single context, then this puts pressure on evaluators to accommodate two different and prima facie incompatible ways of fixing contextual standards. While it is predictably unobvious how evaluators should and do respond in such cases, there are many options that are worth taking seriously, some of which do and some of which don't allow us to construe such cases as involving disagreement. (For discussion of similar cases with respect to contextualism about knowledge, and examination of some of the options, see DeRose (2009, ch. 4).)




Figure 1: Adelson's checkers shadow illusion.

*not* exemplify the very same color. And he complains that this treatment is counterintuitively revisionary.

Now, as Tye recognizes, the relationalist does have the wherewithal to describe this and similar cases of illusion as involving representational error. Namely, she can say describe the case as involving an error in the attribution of a color that the chips *would* be manifesting were they presented under the same viewing condition. That is, she will indeed say that (in addition to representing *A* and *B* as occurrently manifesting different colors) subjects represent *A* and *B* as being such as to mismatch perceptually were they (as they are in fact not) presented under a shared circumstance — in effect, their visual systems are predicting that the chips would look different (would not match) were they presented in identical circumstances. But, of course, this last representation is erroneous: it turns out that the regions would indeed be a perceptual match were they presented in identical circumstances. Thus, the relationalist has the means to say that, in this and other cases we normally describe as color illusions, the visual system is indeed misrepresenting a color property.

However, Tye objects that this proposal doesn't fit the way we ordinarily think of illusion. For, he writes,

In an illusion, ... there is some feature *x* looks to have and *x* lacks that very feature. There simply is no illusion if there is some feature *x* looks to have and *x* lacks not that feature but some other. ... On Cohen's view, the feature *A* and *B* look to have is that of being different in color for normal perceivers in the given *variable* lighting conditions; the feature they lack is another feature, namely, being different in color for normal perceivers in the very same lighting conditions (5).



Figure 2: Canonical instance of color constancy.

Consequently, he concludes, the description of the case available to the relationalist won't count the case as a color illusion.

But Tye's constraint on descriptions of illusion is unmotivated, and one that relationalism satisfies in any case. On the first point, I agree that we should prefer a theory that enables us to describe what are ordinarily construed as illusions as involving some kind of representational error (something that even Tye concedes relationalism can do). Beyond that, it's hard to see why we should regard the classification of such cases as illusions — more particularly, as illusions under Tye's preferred theoretical description of what illusion amounts to — as a non-negotiable condition of adequacy on a theory of color. On the second point, even if we do treat that condition as non-negotiable, it is a condition relationalists can meet. For there *is* a feature that the relationalist will say *A* is represented as having but lacks. Namely, her story about the checkers shadow case is exactly that the visual system represents *A* as exemplifying this property that, it turns out, *A* lacks:

$\lambda(x)$  (were *x* and *B* presented in a common condition type, they would not be perceptual matches).

It would seem, then, that relationalism supplies description of color illusions that is adequate to both the data and Tye's (unmotivated) demand. While Tye is, of course, free to prefer a different description of the cases, he owes us a reason for thinking that his alternative is superior. As things now stand, I don't see that his complaints about illusion have revealed any shortcoming of color relationalism.

I turn now to color constancy, a canonical instance of which is depicted in figure 2, in which a materially uniform coffee cup is partially illuminated by direct sunlight and partially shaded (there is a luminance edge on its facing surface). Part of what makes this and other garden variety instances of color constancy interesting is that it seems that vision represents two different —



and, in some ill-understood sense, conflicting — things about the items of interest (here the two regions of the cup that lie on opposite sides of the luminance edge). In this instance, for example, there is surely some good respect in which the differently illuminated regions of the cup are represented as being alike (hence the label *constancy*) and also some good respect in which the differently illuminated regions of the cup are represented as being different (hence the ease of discriminating the two). The question that concerns us is how we should describe the relationship between the colors of the two regions in a way that allows us to recognize both the represented respects of similarity and difference.

Now, as Tye recognizes, the color relationalist has such a description to offer. Namely, the relationalist will say (*RR*, 53–57) that the two regions are represented as different in the colors they occurrently manifest, and that they are represented as similar in the colors they would manifest were they presented under a common illuminant. As before, however, Tye finds this description unsatisfying:

...if ...the perceptually distinguishable regions ...manifest different colors, then, on Cohen's account of color, they actually *look* different colors. According to Cohen, then, there isn't color constancy (in the relevant sense). This seems wrong to me and to miss the point. I take color constancy for the purposes of this objection to be constancy in how things *look* color-wise through different lighting conditions. It isn't constancy, period. Cohen fixes up something that gets the latter but he doesn't get the former.

Tye is correct that the relationalist description of the case does not sustain his claim that the two regions are alike in the colors they occurrently look to have. But why suppose that an adequate description of color constancy *must* respect that claim? It is uncontroversial that the claim at issue is not mandated by data about psychophysical matching, since all sides agree that the relationalist description of the case is materially adequate to the data. Nor, so far as I can see, is the claim required by anything else we know about the case. Given this, and so long as the relationalist can capture the idea that there is a sense in which the regions are represented as alike — again, something all sides concede — it's hard to see why the similarity that comes out in the matching data must be captured in the specific theoretical way that Tye prefers. In effect, Tye is complaining that the relationalist description of color constancy is not identical to the theoretical description of it that he prefers. That is true, but not, by itself, in any way damaging to relationalism.

## Regress and Color Experience

Tye's final complaint concerns regress. On a very broad class of views — including but by no means limited to the widely popular representationalist/intentionalist views associated with Harman (1990); Dretske (1995);

Tye (1995), color experiences are constituted by relations to color properties. Following Pautz (2006), let us call such views *standard property relational* accounts of color experience.<sup>8</sup> Tye worries that, in so far as color relationalism understands color properties in terms of relations to color experiences, the combination of color relationalism with any standard property relational view of experience will result in circularity.<sup>9</sup> He takes this to show that, contrary to *RR*'s claimed agnosticism about color experience, relationalists about color can't accept any standard property relational account of color experiences — something he regards as “a high cost to pay for color relationalism” (8).<sup>10</sup>

First a small point about the accounting. Readers will have to judge for themselves whether the price of giving up standard property relational theories of experience is objectionably high. But it is worth reminding ourselves that there are a wide variety of alternative accounts of color experience from which to choose, including neural state type identity theories (Hill, 1991), adverbial theories (Sellars, 1975), sense-datum type views (Jackson, 1977; Peacocke, 1984), and primitive phenomenal type views (Chalmers, 1996). The present point is not that the color relationalist should endorse any one of those theories in particular. It is that if, indeed, color relationalism could not be coherently combined with standard property relational accounts of experience, this would by no means leave the color relationalist without theoretical options.

My second, and more important, point is that the claimed regress doesn't get going without a controversial substitution principle that we have reason to reject. To see this, suppose we accept both:

**Color relationalism:** colors are constituted partly by relations to color experiences, and

**Standard property relationalism about color experience:** color experiences are constituted partly by relations (viz., relations of looking/appearing) to colors.

The combination of these views will license the adoption of identity claims of the form

⌈Red = looks/appears  $\langle p_1, \dots, \text{Red}, \dots, p_n \rangle$ ⌋.

<sup>8</sup>Pautz's terminology is “standard relationalism”. Both his and my terminology are intended to make room for non-standard relational views on which color experiences are constituted by relations to properties distinct from but corresponding to the colors — e.g., primed colors (Peacocke, 1984), perfect colors (Chalmers, 2006), etc.

<sup>9</sup>Strictly speaking this concern needn't apply to every form of color relationalism; as mentioned in *RR*, 169, there are versions of color relationalism according to which the color-constitutive relations to subjects do not themselves make reference to color experiences or in any other way involve relations to color properties. Such forms of color relationalism could, of course, be combined with standard property relational accounts of color experiences without threat of regress. I'll put this subtlety aside.

<sup>10</sup>Though I respond to this and related worries in *RR*, §§6.4–6.5, Tye doesn't say just what he thinks goes wrong in my attempts to defend relationalism on these matters. In what follows I'll attempt to set out my views more clearly in the hope of assuaging his dissatisfaction on this score.

That is, the combination would license identity statements on whose left hand sides are the names of colors, and on whose right hand sides are looks/appears relations to colors. But is it true that the occurrence of color names on both left and right hand sides of such identity statement by itself generates a regress? Crucially, no. The regress won't arise without the assumption of some principle to license the substitution of their entire right-hand sides for the occurrences of the color terms that those right-hand sides embed.<sup>11</sup>

Yet there is good reason, independent of any issues about color or color relationalism, for rejecting the contemplated substitutions. Namely, the contexts within which we are considering making these substitutions are, on their faces, non-extensional: quite generally speaking, substitution of extensionally equivalent expressions within these contexts appears not to preserve truth. Thus *a* can look/appear Morning-Starish without looking/appearing Evening-Starish even if being Morning-Starish is identical to being Evening-Starish; *a* can look/appear wet without looking/appearing covered in the appropriate way by H<sub>2</sub>O even if being wet is identical to being covered in the appropriate way by H<sub>2</sub>O; and *a* can look/appear humorous without looking/appearing disposed to cause characteristic amusement reactions in appropriately situated cognizers, even if being humorous is identical with being disposed to cause characteristic amusement reactions in appropriately situated cognizers.

The upshot, then, is that the regress Tye treats as revealing a fatal weakness in color relationalism can get off the ground only if we accept the propriety of substitutions that we have independent reason not to accept. Once we reject the problematic substitutions, the alleged regress — and, with it, the alleged clash between color relationalism and standard property relational theories of color experience — evaporates.

## Relationalism and De Se-ism

I turn now to Egan's concerns about the compatibility and incompatibility of distinct color representation. Egan agrees with me (and against Tye) that relationalism allows for both the kind of compatibility that underwrites ecumenicism about perceptual variation and the kind of incompatibility that explains color disagreement. But he wonders whether we might not prefer an alternative story (one he stops short of endorsing fully) for securing these competing desiderata.

On this alternative de se-ist view, colors are not relational properties, but "centering features" — viz., functions from centered worlds to extensions. Thus, on this view, a (visual/cognitive/linguistic) representation that attributes green to *a* delivers the centered worlds proposition that is true in

---

<sup>11</sup>Needless to say, there may be further reasons — over and above fears of regress — for dissatisfaction with the sorts of identity claims at issue, and possibly even for taking these dissatisfactions as impugning color relationalism. I attempt to defend relationalism against at least some of these considerations in *RR*, ch. 6.

$\langle w, t, i \rangle$  iff  $a$  is disposed to look green to  $i$  in the circumstances  $i$  occupies at  $t$  in  $w$ .

Among its many other virtues, de se-ism is designed to secure the possibility of both compatibility and incompatibility between distinct color representations.

On the one hand, the de se-ist explains incompatibility this way. Suppose  $S_1$  attributes *unique green* to  $a$ , while  $S_2$  attributes *bluish green* to  $a$ . For the de se-ist, this means  $S_1$ 's representation delivers the centered worlds proposition that is true in  $\langle w, t, i \rangle$  iff  $a$  is disposed to look unique green to  $i$  in the circumstances  $i$  occupies at  $t$  in  $w$ ; and  $S_2$ 's representation delivers the centered worlds proposition that is true in  $\langle w, t, i \rangle$  iff  $a$  is disposed to look bluish green to  $i$  in the circumstances  $i$  occupies at  $t$  in  $w$ . But the truth conditions of  $S_1$ 's and  $S_2$ 's representations are disjoint: there's no centered world such that the extension *unique green* determines relative to that center overlaps with the extension *bluish green* determines relative to the very same center. Or, in other words, there are no individuals disposed to look both unique green to  $i$  in the circumstances  $i$  occupies at  $t$  in  $w$  and bluish green to  $i$  in the circumstances  $i$  occupies at  $t$  in  $w$ . (Correspondingly, the de se-ist can explain color *agreement* as an exact match in the extensions two attributions determine relative to a fixed center.)

On the other hand, de se-ism also provides a way in which distinct color representations can be compatible in the sense needed to secure ecumenicism about perceptual variation. Namely, even if  $S_1$ 's and  $S_2$ 's representations have disjoint truth conditions relative to a fixed center, they can both be true relative to distinct centers. So when  $S_1$  attributes unique green relative to  $\langle w, t, S_1 \rangle$ , and  $S_2$  attributes bluish green relative to  $\langle w, t, S_2 \rangle$ , these attributions can both be veridical.<sup>12</sup>

I hope it is clear that de-seism and relationalism are closely related. Crucially, both views agree that some kind of relative/relational subject-involving element is needed to accommodate the observed range of perceptual variation, though they disagree in where they locate that element. The

---

<sup>12</sup>A related difference between the two views concerns property proliferation. The key to my version of ecumenicism lies in the proliferation of distinct relational properties, all of which can be (compatibly) exemplified by a single object. Egan's implementation of ecumenicism has no need to proliferate properties, but instead proliferates centers with respect to which we can evaluate the extension of each property. I'm not sure in the end there's much of a substantive issue here.

I do, however, want to forestall the impression (not one I think Egan holds) that the relational property proliferating view amounts to some hideous violation of ontological parsimony. Perhaps it is objectionably unparsimonious to proliferate non-relational properties (though I take no official stand). Even so, 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. Likewise, I suggest, for a theory of color that is already committed to the existence of objects, subjects, and viewing conditions (cf. *RR*, 133–136).

relationalist locates the subject-involving element in the properties (by making the properties relational), while the de se-ist is agnostic about the metaphysics of the properties themselves, and instead locates the subject-involving element in the way that the properties map different centered worlds into possibly different extensions.

It seems to me that the de se-ist view is powerful, interesting, and almost surely correct about *some* properties (e.g., it is utterly compelling as a story about the property shared by those thinkers who believe their own pants are on fire). My main question about the view is how far it should be extended, and, specifically, whether it should be extended to color properties.

One way to frame this question is to note that the de se-ist proposal is in principle no less applicable to properties that are paradigmatically and uncontroversially relational, such as *moving at 60mph*. On the de se-ist analysis, we could say that this paradigmatically relational property is not (or may not be), after all constituted in terms of a relation to a subject. Instead, the view would treat *moving at 60mph* as a function from centered worlds to propositions, so that the truth of an attribution of this property to an individual depends on the triple  $\langle w, t, i \rangle$  relative to which we evaluate.

The exercise of applying the de se-ist strategy to such cases invites a couple of concerns.

My first concern about the contemplated application of the de se-ist strategy is not about the propriety of the descriptions it allows, but instead about its comparative utility. For, supposing that hard-won inquiry has revealed the need for a subject-involving element in understanding a certain property or family of properties (as Egan and I agree it has with color properties) it strikes me as odd to prefer a metaphysics on which the 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 to me more useful to mark the hard-won knowledge by encoding the property's relationship to subjects directly in our metaphysical account of its nature. De se-ism 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.

My second, related, concern is about systematization. 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 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 de se-ism'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 de se-ist can model the subject-involving extension selections that color properties make, there is nothing in the de se-ist metaphysics of color to explain why all of the different colors are appropriately modeled in that very same way. This is one respect in which, it seems to me, relationalism is preferable to de se-ism.<sup>13</sup>

So, given the close affinities between the two positions, should de se-ism be counted a species of relationalism? I agree with Egan that there are principled ways of answering this question either affirmatively or negatively, depending on just which criteria one takes to be essential to relationalism. However, the letter of the law (as handed down in *RR*) says that relationalism is the view that colors are constituted in terms of relations between subjects and objects (*RR*, 8), and de se-ism doesn't meet this condition: it denies that colors are constituted in terms of such subject-involving relations (and instead locates the subject-involving element it recognizes in the way in which colors determine extensions). My official answer, then, is that de se-ism is not a form of relationalism, but instead a clearly important and interesting close cousin.<sup>14</sup>

## References

- Byrne, A. and Hilbert, D. R. (2003). Color realism and color science. *Behavioral and Brain Sciences*, **26**(1), 3–64.
- Byrne, A. and Tye, M. (2006). Qualia ain't in the head. *Noûs*, **40**(2), 241–255.
- Chalmers, D. (1996). *The Conscious Mind: In Search of a Fundamental Theory*. Oxford University Press, New York.
- Chalmers, D. (2006). Perception and the fall from Eden. In T. S. Gendler and J. Hawthorne, editors, *Perceptual Experience*, pages 49–125. Oxford University Press, New York.
- Cohen, J. (2012). Précis of *The Red and the Real*. *Analytic Philosophy*.
- DeRose, K. (2009). *The Case for Contextualism*. Oxford University Press, Oxford.
- Dretske, F. (1995). *Naturalizing the Mind*. MIT Press, Cambridge, Massachusetts. Originally delivered as the 1994 Jean Nicod Lectures.

---

<sup>13</sup>Another possibility that suggests itself is that one might somehow combine a relationalist metaphysics of color properties with a de se-ist semantics of color predicates. The hope would be that the relationalism might provide a principled and general account of *why* color predicates are centering features, as the de se-ist proposes. I'm not certain whether any such hybrid position could be worked out and motivated in an acceptable way, but it seems worthy of consideration.

<sup>14</sup>Thanks to David Brink, Craig Callender, Adam Pautz, and Chris Wüthrich for helpful comments on earlier drafts.

- Harman, G. (1990). The intrinsic quality of experience. In J. Tomberlin, editor, *Philosophical Perspectives: Action Theory and Philosophy of Mind*, volume 4, pages 31–52. Ridgeview Publishing Company, Atascadero, California.
- Hill, C. S. (1991). *Sensations*. Cambridge University Press, Cambridge.
- Jackson, F. (1977). *Perception: A Representative Theory*. Cambridge University Press, New York.
- Pautz, A. (2006). Can the physicalist explain colour structure in terms of colour experience? *Australasian Journal of Philosophy*, **84**(4), 535–565.
- Pautz, A. (2010). Review of Jonathan Cohen, *The Red and the Real*. *Notre Dame Philosophical Reviews*.
- Peacocke, C. (1984). Colour concepts and colour experiences. *Synthese*, **58**(3), 365–381. Reprinted in Rosenthal (1991), 408–416.
- Rosenthal, D. (1991). *The Nature of Mind*. Oxford University Press, New York.
- Sellars, W. (1975). The adverbial theory of the objects of perception. *Metaphilosophy*, **6**, 144–160.
- Tye, M. (1995). *Ten Problems of Consciousness: A Representational Theory of the Phenomenal Mind*. MIT Press, Cambridge, Massachusetts.
- Tye, M. (2012). Cohen on color relationalism. *Analytic Philosophy*.