

Chromatic layering and color relationalism

Jonathan Cohen*

Herzog was all too well aware of the layers upon layers of reality — loathsomeness, arrogance, deceit, and then — God help us all! — truth, as well.
— Saul Bellow, *Herzog*.

Brown (2015) highlights cases of “chromatic layering” — scenarios in which one perceives an opaque object through a transparent volume/film/filter with a chromatic or achromatic content of its own — as a way of reining in the argument from perceptual variation I’ve used to motivate a relationalist account of color properties (Cohen 2004, 2009). That argument (about which more anon) generalizes widely across a very wide range of circumstantial and perceptual parameters which, when they take on different values, lead to a difference in the experienced color of a fixed stimulus. Such parameters include, e.g., viewing distance, viewing angle, retinal photoreceptor type ratios, chromatic surround, and so on. The intended conclusion of the argument is that color properties are best understood in terms of relations to the relevant parameters. But Brown urges that the argument in question does not generalize smoothly to all types of perceptual variation — in particular, that it fits poorly in layering cases in which there is either experiential fusion or scission. While he doesn’t want to reject the relationalist description of variation in all cases, he does suggest that there is an alternative, and preferable, description available for layering cases, and that the availability of this alternative puts limits on the scope of both relationalist arguments and relationalist conclusions.

In what follows I want to suggest that the obstacles for the argument from perceptual variation (hence for relationalism) posed by cases of layering are merely apparent. I argue that relationalism can be extended smoothly to cases of layering, and indeed that denying this extension, as Brown proposes, contravenes ordinary and more or less universally accepted canons of rational inquiry. I’ll begin by reviewing the argument from perceptual variation, and showing how it motivates color relationalism (§1). Next, I’ll present Brown’s layering cases, and explain how he thinks they pose challenges to the argument from perceptual variation (§2). Then I’ll attempt to defuse these challenges: I’ll argue that there are no grounds for regarding either fusion or scission as problems for the argument from perceptual variation or color relationalism (§§3–4). Finally, I’ll reply to Brown’s methodological objections; I’ll show that, contrary to his accusations, the relationalist treatment of layering rests on (and Brown’s preferred analysis depends on flouting) a very weak and well-motivated conception of uniformity that we all have reason to accept (§5). I’ll conclude that the challenge from layering is a challenge relationalists can meet.

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

1 Color relationalism and the argument from perceptual variation

In earlier work (esp. [Cohen 2004, 2009](#)) I have argued for color relationalism — the view that objects' colors are constituted in terms of relations those objects bear to perceiving subjects and perceptual circumstances. The thought is that, just as things count as instances of *being a teacher* or *being humorous* not by virtue of their intrinsic makeups, but by standing in the right (viz., pedagogical/amusing) sort of relation to relata (viz., students/appropriately equipped cognizers), so, too, things count as exemplifying colors by standing in the right relation to perceivers and perceptual circumstances.

My case for relationalism revolves around a traditional, non-deductive, empirically motivated argument based on perceptual variation.¹ Roughly this non-deductive argument begins with the empirical claim that there is significant inter- and intra-personal variation in perceptual 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 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 this view gives us the best way of implementing the desired ecumenicism.

Though the argument form just given generalizes widely, it is perhaps easiest to appreciate its force by considering its application to a single, simple instance of perceptual variation, such as that involving simultaneous contrast in a traditional center-surround configuration such as figure 1. Though the two central patches in



Figure 1: A traditional center-surround configuration. The two central patches are intrinsically qualitatively identical, but the patch set against a less yellow /more blue surround on the left appears more yellow /less blue than the patch set against a less blue /more yellow surround on the right.

this configuration are qualitatively identical in their non-relational properties, they

¹Since [Brown \(2015\)](#) discusses this argument, I give only a compressed characterization here. For a fuller presentation and much more discussion and defense, see [Cohen \(2009\)](#).

look different (in color) depending on the surround against which they are placed. Assuming, standardly, that the way the patches look represents their colors, it follows that the visual system represents the colors of the patches differently as a function of the surround. If so, we can ask: which, if any of the (psychophysically distinguishable) representations of the patches is veridical?² The logically possible answers are: neither, the first to the exclusion of the second, the second to the exclusion of the first, or both. I suggest that the first option (neither right) is unacceptably skeptical/revisionary, and that the second and third answers (one representation exclusively) are unacceptably arbitrary — that it is hard to imagine a well-motivated, principled, and non-question begging reason to believe that either representation is uniquely veridical. Assuming we want to avoid both revisionary skepticism and *ad hoc* stipulation when possible, this suggests that, *ceteris paribus*, we should prefer the view that both representations are veridical. My further suggestion is that we should endorse color relationalism because it gives the best way of understanding how this answer could be true. For relationalism allows us to say that the patch has two compatible colors: it is (simultaneously, all over) light yellowish green with respect to the first perceptual circumstance, and a slightly darker bluish green with respect to the second.³ Color relationalism is attractive because it allows us to avoid a hard choice between unpalatable alternatives.

2 Layering: scission and fusion

Brown is interested in the challenges posed to the argument from perceptual variation (and therefore to color relationalism) by instances of chromatic layering. He centers his case against the argument on two types of scenarios: cases of “scission,” in which visual experience separates out multiple objects located along the same line of sight and assigns colors to each of them, and cases of “fusion,” in which experience fails to distinguish between multiple objects along the line of sight, and consequently fails to attribute separate colors to each of them. I’ll follow Brown in discussing these in turn.

2.1 Scission: Variation without relationalism?

We can begin with the scission case (what he calls “the Layering Scenario”) represented in figure 2, where the perceiver perceives an opaque object O first through a transparent filter F_1 (figure 2a) and then through a second transparent filter F_2 . (figure 2b).

Now, this sort of case may initially appear to be grist for the relationalist’s mill — yet another instance of perceptual variation to which we can apply the argument from variation, as outlined above. After all, just as in the case of the center-surround configuration depicted in figure 1, and used as a paradigm for explaining the argument from perceptual variation, here we have a case in which the subject undergoes psychophysically distinguishable experiences of the color of the object O , each of which purports to represent O ’s color, each of which can be assessed for veridicality, and between which there is no obvious way to make a principled and exclusive choice.

However, Brown urges that we should not assimilate the layering scenario to standard cases of perceptual variation in this way. He accepts that, in the layering

²The question here is metaphysical, not epistemic: it is not ‘how do we know which of the perceptual effects veridically represents the patch’s color?’, but ‘what makes it the case that one of the perceptual effects (as opposed to others) veridically represents the patch’s color?’.

³Analogy: a single individual can be a teacher of you, and, simultaneously, not a teacher of me.




Figure 2: (a) Perceiver views an opaque object O through a transparent filter F_1 , and then (b) through a different transparent filter F_2 .

scenario, there is a difference in the total color experience involving one target object (“ O_{white} ”) and a numerically distinct transparent object (“ T_{grey} ”) along the line of sight between the observer and the target object.⁴ But, Brown advises, we shouldn’t take this to indicate that the cross-contextually varying experience of the subject is varying in what color it attributes to the single target object, as needed by the argument from perceptual variation:

Clearly O_{white} is causally involved in bringing about the experience ... O_{white} -through- T_{grey} , and in that sense O_{white} has a role in generating these distinct colour experiences in the [Layering] Scenario. But it would be erroneous to hold that O_{white} itself generates these distinct colour experiences, or to hold that these distinct experiences are simply experiences of the colour of O_{white} In order for [the argument from perceptual variation] to create a rationale for Relationalism, the variable colour experiences drawn from the relativity of colour must be experiences of the target object’s colour (i.e., “of a single color stimulus”). But in our Layering Scenario the difference between the experiences O_{white} -through- T_{grey} and O_{white} -through- T_{yellow} is in a fundamental sense not a difference in the

⁴Brown’s labels for objects (‘ O_{white} ’ for the opaque object terminating the line of sight, ‘ T_{grey} ’/‘ T_{yellow} ’ for transparent filters between the perceiver and target object), which he uses to pick out these objects across variations in perceptual circumstances, encourage the assumption that these objects have stable colors that are independent of (and persist across variations in) perceptual circumstances. Since that assumption is crucially under dispute, I have chosen more neutral labels for the objects (‘ O ’, ‘ F_1 ’, ‘ F_2 ’) in my own description of the case, but will leave Brown’s nomenclature as it stands when quoting him.

experienced colour of O_{white} , it is a difference in the experienced colour of something else along the LOS [=line of sight] on which O_{white} is currently located. Thus while, as it were, the total colour experience along this LOS changes, and O_{white} is the terminal visible object along that LOS, the experienced colour of O_{white} is stable, and there is no motive to infer that O_{white} 's colour is relational (186).

According to Brown, then, since the successive and psychophysically distinguishable experiences in question (call them e_1 and e_2) do not attribute their chromatic content entirely to the very same individual, O , they should not be construed (as in standard instances of perceptual variation) as putatively rival representations of O 's color. True, e_1 and e_2 differ in their overall chromatic content; but that is because of the differential contributions made to the overall chromatic contents of those experiences by objects *other* than O — namely, F_1 and F_2 . As far as O itself, e_1 and e_2 agree in the color they attribute to O (“the experienced colour of O_{white} is stable”). As such, there is no reason for thinking we face any hard choice of which of the two veridically represents O 's color. They agree, so we can unproblematically regard both as veridical without any need for treating the properties they attribute as relativized or relational.

In short, Brown's contention is that the best description of the psychophysical variation in layering scenarios obviates the hard choices to which relationalism was proposed as an escape, and thereby leaves relationalism unmotivated (at least in these cases).

2.2 A retreat to fusion?

Given the problems just sketched, Brown imagines that relationalists might prefer to run the argument from variation only on cases of fusion. But he objects that in that case the argument from perceptual variation will fail for another reason.

After all, the argument from variation crucially turns on the premise that there is symmetry between perceptual variants — that there is no principled, non-arbitrary, non-question-begging reason for treating just one of the variants as uniquely veridical. But Brown worries that this symmetry premise is plainly false in fusion cases. In his view, chromatic fusion is chromatic confusion. And, as such, the standard epistemic value of avoiding informational conflation counts as a principled, non-arbitrary, non-question-begging reason to prefer over fused variants either (i) a layering/non-fused experience that separates out the contribution of O from other objects on the line of sight, or else (ii) variants in which there is no other object making a contribution to our chromatic experience:

If one's aim is to discern the colour of the book [O], then a colour experience in which the book's contribution to colour is distinctly experienced should be preferred over one in which the book's [O 's] and cellophane's [F_1 's] contributions are fused together. Fusion experiences yield information, not about the book's colour simpliciter, but about a relation between the book's colour and the cellophane's colour. This relational information is valuable, but if one is seeking information about the book's colour, the information is arguably inferior to information in which the book's contribution to colour experience is distinctly packaged, or separated out from the contributions of other objects... If one is seeking to learn of the book's colour, it is reasonable to privilege the former over the latter. Privilege should also be

granted to experiences in which only the book is contributing to the colour experience along a [line of sight] (187).

The upshot, then, is that the argument from perceptual variation fails for fusion cases, so that restricting its application to such cases won't save the argument, hence won't allow it to do its work in motivating color relationalism.

3 Fusion revisited

Before I respond to Brown's (as I read him, more central) criticisms involving scission layering scenarios in §4, I want to address more briefly his considerations about fusion cases.

As we have seen, Brown thinks the argument from perceptual variation breaks down in such cases because ordinary epistemic norms break the symmetry between perceptual variants needed by the argument, and breaks it in favor of variants that avoid informational conflation. But there are two reasons for thinking that this point is ineffective against the argument from perceptual variation.

First, though the objection turns on thinking of fusion as a kind of informational conflation that is in conflict with epistemic norms, that characterization is unsupported and close to question-begging. It is true that the fusion experience yields information about O in relation with F_1 , rather than about O simpliciter. However, there's no reason to assume that that the former information is less epistemically valuable than the latter — not even when one is "seeking information about the book's colour." That conclusion only follows if there is some sort of standing epistemic preference for experientially representing objects via their non-relational features. But why accept that? In the envisaged case, O and F_1 do have a joint causal effect on the perceiver, and the fusion experience under consideration (veridically) represents O in its relation to F_1 . That content is information about the book and its color. To insist that the content of a fusion experience is epistemically inferior to that of a non-fused alternative experience is just stipulative, and certainly won't be accepted by Brown's relationalist opponent.⁵

Second, we should recall just what sort of symmetry is at issue. The targeted symmetry premise of the argument from perceptual variation is the claim that there is no well-motivated, principled, and non-question begging reason for *singling out one representation as uniquely veridical*. But even if (contrary to what I just said) experience did come with an epistemic preference for non-relational feature representations, that wouldn't by itself break the symmetry on which the argument from perceptual variation is grounded. It wouldn't be a reason for believing one of the variants is veridical to the exclusion of others. After all, even Brown's own explicit description of the case has it that the fused color experience of O -cum- F_1 is no less veridical than a non-fused color experience of O . Thus, even granting Brown's assumptions about the epistemic norms governing experience does nothing to undercut the symmetry premise on which the argument from perceptual variation depends.

⁵Compare: It would be absurd to demand an account of whether someone is, say, a sister, or a teacher, in terms that don't mention ("conflate the contributions of") other persons. Likewise, it seems inappropriate to insist on applying some putative epistemic preference for "non-conflated" descriptions of our chromatic experience before we have reached a verdict about whether the contents of chromatic experiences are relational. Saying this much is not to beg the question in the opposite way by assuming a contrary epistemic norm; rather, it is to suggest neither sort of epistemic norm should be invoked at this stage in the dialectic.

In short, I don't see that what Brown has said about fusion cases is damaging to the relationalist's attempt to motivate her view by appeal to the argument from perceptual variation.

4 Layering for relationalists

But I take it that it is the case of scission, rather than fusion, that lies at the heart of Brown's criticisms of the argument from perceptual variation and color relationalism. I now want to argue that cases of scission (henceforth, cases of layering) are similarly ineffective in limiting the scope or force of the case for color relationalism.

The first point to make is that (as Brown is aware) relationalism *does* make room for a description of layering cases. We have already seen that, for relationalists, object colors are relational properties containing a parameter for perceptual circumstances, which they think must be specified in terms of whatever features of the circumstance make a psychophysically identifiable difference to experience (but where the variants have an equal claim to veridicality). Viewing an opaque object O through a transparent filter, as in the cases depicted in figure 2, certainly appears to make a psychophysically identifiable difference of this sort to our experience of O . So it is entirely consistent with normal relationalist methodology to describe scission cases by the same means. A relationalist will say that the relational color O manifests when viewed through filter F_1 is, as it might be, *grey to S in perceptual circumstance C_1* , where C_1 involves (among other possibly relevant parameters) the presence of F_1 along the line of sight. And she will claim that the color O manifests when viewed through filter F_2 is, as it might be, *yellow to S in perceptual circumstance C_2* , where C_2 involves (among other possibly relevant parameters) the presence of F_2 along the line of sight. As usual, this maneuver will allow the relationalist to hold that both of the experiences veridically represent a (relational) color that the cup has (all over, simultaneously).

As I say, Brown is aware that the relationalist has this story to tell about layering cases; I take it this is why he denies explicitly (183) that they are straightforward counterexamples to relationalism. Rather, his complaint is that the available relationalist account is less than satisfactory, hence that it would be "premature" (183) to accept it. His stated reason for dissatisfaction with the relationalist's alternative description of layering cases is that it wrongly assigns chromatic content to just the one perceptual object O rather than to the two distinct objects O and F_1 ($/F_2$), as he prefers (186).

But that's not quite right. As we have just seen, the relationalist thinks of the case involving filter F_1 as one in which O manifests the color *grey to S in perceptual circumstance C_1* , where C_1 crucially involves (inter alia) the presence of F_1 along the line of sight. True, this is an assignment of color to O : it is an ascription of a color that O , rather than F_1 , manifests. But that color constitutively involves the object F_1 . So it is not true that the relationalist's description of the chromatic content in the scene mentions only the single object O ; rather, it mentions O in a circumstance involving the numerically distinct object F_1 . Nor is it true that the relationalist can only assign color to one of these two objects; for, over and above the relational color she attributes to O , she'll also want to allow that object F_1 exemplifies a distinct relational color constituted in terms of a relation to the perceptual circumstance in which it is perceived (in this case that circumstance will involve the simultaneous perception of O along the same line of sight; see §4.2.)

As such, the choice between descriptions of layering cases we face is not a choice between descriptions that recognize one as opposed to two objects, and it is not a choice

between descriptions that recognize one as opposed to two colors. Rather, it is a choice between a description that treats O 's chromatic content as independent of F_1 and one that treats O 's chromatic content in terms of its relation to F_1 . And now our question is: why should we prefer a description of the first type, as Brown prefers, over a description of the second type, as the relationalist prefers?

Though Brown is less fully explicit on this point than one might have hoped, his paper suggests a few possible, and somewhat overlapping, answers.⁶

4.1 Obviously wrong?

In some places Brown suggests that it is just obvious his own description of the case is the uniquely correct one. For example, at one point he considers and rejects the relationalist redescription of O 's in terms of its circumstance, so that its color changes when seen through different filters F_1 and F_2 :

... aren't you just experiencing the opaque thing to be differently coloured across these contexts? This is not a friendly amendment, but instead an attempt to eliminate experienced colour layering. There is no such simplification: experienced colour layering can only be adequately described via a complex such as O_x -through- T_y (185).

And later on the same page he asserts that

we may additionally suppose that the object colours are constitutively non-relational, that our three objects are intrinsically white, grey and yellow, respectively (185).

But this is unpersuasive. Brown's insinuations that the phenomenon "can only be adequately described" in the way he favors, and that the correctness of the non-relational description is something "we may additionally suppose" are question-begging in the context of our search for a reason to prefer the circumstance-independent description he favors over the circumstance-involving alternative description favored by the relationalist.

4.2 Collapses onto fusion?

A second possible answer, perhaps offered in connection with the first (perhaps this is how we should read Brown's characterization of the relationalist description as "an attempt to eliminate experienced colour layering") involves the worry that the relationalist description wrongly reduces all instances of scission to cases of fusion.⁷ But this charge (whether or not Brown would make it) is inaccurate. The relationalist will hold that scission cases differ from fusion cases because, in the former but not the latter, experience attributes a further relational color (additional to the relational color of O) to the filter F_1 .

To see the point, notice that a relationalist has no reason to reject the idea that visual experience at a time can attribute colors to more than one object in a single stimulus

⁶In attempting to consider a broad range of possible motivations, in several places below I'll go beyond claims to which Brown is explicitly committed. My critical reactions are aimed at the proposed motivations themselves, whether or not Brown would endorse them.

⁷I'm genuinely unsure whether this worry is one that Brown has or takes seriously. Note that, though the quoted phrase from p. 185 suggests that he thinks of the relationalist's description as ultimately inadequate to the phenomenon of layering, in other places (e.g., 179, 183) he seems much more open to that description.

configuration. After all, that’s exactly what she will say about a center-surround configuration like the one depicted in figure 1.⁸ She’ll hold that visual experience attributes two relational colors in this case. First, she’ll say, it attributes to the central patch a relational color involving the presence of the surround as part of the perceptual circumstance; second, she’ll say, it additionally attributes to the surround a relational color involving the presence of the central patch as part of the perceptual circumstance. Moreover, this approach can be carried over straightforwardly to cases of scission such as that depicted in figure 2a. First, the relationalist will say of such a case, experience attributes to O a relational color involving the presence of F_1 as part of the perceptual circumstance; second, she’ll say, it attributes to F_1 a relational color involving the presence of O as part of the perceptual circumstance.

But now we are in a position to see that this consistent extension of ideas already available to relationalists (and, indeed, already exploited in their description of center-surround variation) provides a simple way of distinguishing cases of scission from cases of fusion. A relationalist will say that fusion cases are those in which experience attributes just one relational color to an object along the line of sight (here O), while in fission cases experience attributes numerically distinct relational colors to multiple objects along the line of sight (here O and F_1).⁹

4.3 Neglects constancy?

A third possible reason for favoring Brown’s preferred description of layering is connected with his reasonable suggestion (189–190) that there is a kind of color constancy associated with our experiences of O through different layering filters such as F_1 and F_2 . As he says, there is plausibly some respect of chromatic sameness between two such experiences of O , and we should want a conception of the cases that makes room for this fact. One might fear that the relationalist descriptions of the two layering cases (e_1 represents O as grey to S in a condition involving the mediation of F_1 ; e_2 represents O as yellow to S in a condition involving the mediation of F_2) unacceptably makes no room for the (as yet unanalyzed) chromatic sameness between the cases.

⁸ Brown appears to accept the relationalist description of ordinary center-surround perceptual variation (“the prevalence of rich phenomenon like center-surround contrast effects should tempt us all to be Relationalists of some sort,” 191; cf. 184). To be fair, in other places he endorses the relationalist treatment of these cases more provisionally (“For simplicity I will assume that the lesson to draw from center-surround contrast effects is that some variant(s) of Relationalism are true,” 182; cf. 184–185). Still whatever his ultimate views on the matter, it seems fair to regard Brown as accepting the relationalist description of center-surround variation — at least *arguendo*, and perhaps more fully — in assessing his challenge from considerations about layering.

⁹Note that this description of fusion has a parallel in traditional center-surround cases as well — in such cases, typically labeled ‘color induction,’ the surround influences the chromatic experience of the central patch but is not assigned a chromatic content of its own.

Indeed, in view of the far-ranging structural parallels between layering and center-surround cases, it’s tempting to think of the former as a special type of center-surround configuration — viz., one where the surround and central patch happen to be located at different depths relative to the perceiver. In principle, the choice between Brown’s favored two-independent-properties-of-two-objects description and the relationalist’s favored one-object-in-a-circumstance-involving-a-second-object description should be available for any of these cases (cf. §5). In this respect, it seems to me that Brown’s challenge is not ultimately about the phenomenon of layering, per se, but about the much more general questions of whether, how, and in what conditions color experience should be described as distributing the chromatic contents of single intentional objects over multiple entities in the environment. It will come as no surprise that I think color relationalism offers a productive and general set of answers to these questions.

Again, this complaint is unpersuasive; for, as Brown points out, there are indeed relationalist resources for accounting for the desired respect of chromatic sameness. Thus, I have proposed in other work (Cohen 2008) that the visual system achieves color constancy by effectively representing, over and above the occurrent colors of perceived items *A* and *B*, counterfactual verdicts about whether *A* and *B* would match if perceived under a shared perceptual circumstance that only one of them actually occupies. On this proposal, the chromatic sameness the visual system represents between e_1 and e_2 is the (true) content that there would be a match between two chromatic experiences generated by the perception of *O* under the same perceptual circumstance — either two experiences of *O* in a circumstance involving mediation by F_1 but not F_2 , or two experiences of *O* in a circumstance involving mediation by F_2 but not F_1 .¹⁰

Brown objects that, while this account may be acceptable in at least some cases, it may be false of others, and that the best methodology requires following the data where they lead:

...one's goal should not be to try to force all instances of colour constancy into one model (be it a layering, counterfactualist or some other model) to accommodate one's broader objectives Our goal should be to examine various instances of colour constancy with various models in mind to try to uncover which one or more of them best explain individual or classes of cases. The potential result is that colour constancy is a varied phenomenon, having some instances of layered constancy, others of counterfactual constancy, and so on (190).

I confess that if there's an objection here, I don't see what it is. Of course it's true that theories should be independently motivated and constrained by the data. In the present case, however, Brown has already accepted (though perhaps only provisionally; see note 8) that relationalism is independently motivated and correct about many cases; moreover, it appears to provide a description of color constancy that is entirely consistent with the data on hand. Unless and until someone identifies some difficulty (empirical or otherwise) with that description, all we can ask from the relationalist (or any other theorist) is that she provide consistent, well-motivated, and empirically adequate accounts of relevant phenomena. So far as I can tell, the relationalist is as far along in meeting this demand as anyone else.

4.4 Forced?

Brown's final complaint against the relationalist description of layering is that (as he complained about the relationalist description of color constancy just reviewed), he finds it forced. Thus, he urges that "one's goal should not be to try to force all instances ... into one model ... to accommodate one's broader objectives" (190), and that "succumbing" to a uniform treatment that applies to both center-surround and layering cases risks making us "blind to surprising possibilities" (191).

¹⁰Some (e.g. Tye 2012) have objected that this "counterfactualist" account of color constancy is unacceptably revisionary in only recognizing a non-occurrent respect of similarity. But since the psychophysical data about color constancy (typically, subjects' matching judgments) are entirely agnostic on this point, surely the onus is on Tye and others who would reject the counterfactualist view to show that the needed respect of similarity must be an occurrent rather than non-occurrent respect. As matters stand, this burden remains unmet.

Once again, I find this objection puzzling. I agree that the relationalist descriptions at issue are motivated by the aim of making relationalism consistent with the data, and that proponents of other views may prefer their own descriptions. But the mere fact that these descriptions are enlisted in the defense of a theoretical commitment — even a prior theoretical commitment — is no strike against them. What would count as a strike against such a description would be a demonstration that it is empirically inadequate or otherwise objectionable; but no such demonstration has been provided.

5 Uniformity

As we have seen, Brown disagrees with relationalists in finding relevant differences between traditional center-surround configurations and layering cases. He accepts that there exist ways of construing both kinds of cases in the very same relationalist-friendly terms; but he thinks that, while such descriptions are appropriate for center-surround cases, they are in various respects unsatisfying when applied to cases of layering. What these considerations suggest to him, then, is that theorizing about the range of cases is best carried out in a case-by-case fashion, rather than by insisting on a uniform treatment:

Here is the general lesson: there are some and potentially numerous instances of experienced colour layering; determining their true extent must be done not by asserting the Uniformity Hypothesis or any such theoretical principle, but by careful examination of relevant cases; due to various confounding factors that examination will not be straightforward. In other words, we cannot assert the Uniformity Hypothesis (190).

Thus stated, it's hard to disagree with Brown's reasonable and attractively modest sentiment: open-mindedness is better than closed-mindedness, and uniformity is at best a defeasible theoretical virtue that can be applied only after ruling out potentially confounding differences. If the "careful examination of relevant cases" that he recommends did indeed reveal such differences, then it would be absurd to insist on treating those cases uniformly. It's just that, as I have argued (§4), the alleged symmetry-breaking disanalogies that Brown offers are illusory. And in this situation, where prior careful examination of relevant cases fails to uncover any reason for treating the cases in different terms, it *is* appropriate to invoke uniformity to choose between descriptive alternatives.

I regard uniformity as a theoretical virtue, not an explicit premise, and surely nothing as strong as Brown's Uniformity Hypothesis ("...that all or most instances of the relativity of colour should be treated as above, namely that each case can be described in a parallel fashion" (183)). I agree with him that, like other theoretical constraints, we should invoke uniformity only after careful consideration of the empirical facts on the ground, and only *ceteris paribus*. But that's just where we now are. We *have* considered the cases carefully and found that *cetera* are indeed *pariba* — that the available broadly empirical, theory-independent, well-motivated constraints don't choose between the descriptive alternatives.

On the contrary, what we know about the cases leaves open the very same theoretical descriptive options with respect to cases of scission and traditional center-surround cases (and, for that matter, many other cases involving variation in the chromatic effects of an object as it moves through different circumstances). On the one hand, we can characterize such cases by saying that the chromatic content experience

attributes to O involves aspects of the perceptual circumstance. Or, on the other hand, we can say that experience parses out its chromatic content, and attributes content to O that is independent of the perceptual circumstance.

There appear, then, to be three options left open after careful examination of relevant cases, i.e., after the recognition that broadly empirical constraints fail to choose between the descriptive alternatives. First, one could accept circumstance-independent color attributions uniformly for both center-surround and layering cases. Second, one could throw uniformity to the wind, and accept different descriptions of each, as Brown suggests.¹¹ Or, third, one could adopt a consistent circumstance-involving description of both cases, as I prefer.

Since Brown and I both reject the first option (we both accept, at least for purposes of the current dispute, the circumstance-involving description of the chromatic content experience attributes at least for center-surround cases), I will put it aside.^{12 13} This leaves the second and third choices on the table.

It's worth being clear that either one of these choices *is* a choice: electing either is making, rather than avoiding, a decision about how to theorize appropriately. But given that we have to make a choice, and that broadly empirical constraints leave open these two options, it's hard to see why we should not at that point invoke the theoretical virtue of uniformity as a reason for preferring the third choice (the consistent circumstance-involving description) over the second (different descriptions of the cases). After all, given that we're in a situation in which there's no compelling, theory-independent reason for treating the (otherwise structurally similar) cases in dissimilar ways, differential treatment of the cases would be arbitrary. It would be unconstrained by reasons in a way that, I take it, we don't want our choices in theory-building to be.¹⁴

My suggestion, then, is that uniformity considerations do give reason for choosing between alternatives left open by other broadly empirical constraints, and that, when

¹¹It's worth adding that, whatever its other virtues and vices, the non-uniform strategy Brown is offering does not avoid commitment to relational colors. Brown is prepared to admit, at least provisionally, that colors are constituted in terms of relations to at least those parameters (such as, in his view, chromatic surround) with respect to which the circumstance-involving description of variation is appropriate. So he's already willing to accept (even if provisionally) that colors are relational with respect to at least some parameters. But if so, then the colors attributed by experience even in scission cases, where Brown is urging us not to adopt a relational description, will in any case be relational with respect to some other parameters. For this reason, the non-uniform strategy on offer will be unattractive to those who (unlike Brown) had hoped to avoid a relationalist conception of color.

¹²Needless to say, therefore, the considerations that follow won't by themselves move someone who, unlike Brown, is not willing to accept, *arguendo*, the relationalist treatment of center-surround cases.

¹³I can't help making the further remark that the uniform circumstance-independent description appears much less plausible when applied to cases of interpersonal perceptual variation — viz., cases where one stimulus O in one perceptual circumstance C varies in perceptual effects on different perceivers. In such cases, it is difficult to describe the chromatic effects as varying as a function of differences in some independent experienced objects. After all, in such cases the variation corresponds with changes in the perceiver/visual system, rather than anything distal. Consequently, extending the circumstance-independent description to such cases would require, implausibly, that experience ascribes chromatic contents to the visual system itself.

¹⁴I take it this point lies at the heart of the attraction of so-called unificationist theories of scientific explanation (see, e.g., Friedman 1974; Kitcher 1989). One doesn't have to be a full-fledged explanatory unificationist to see value in theoretical uniformity (when other things are equal), as I'm suggesting.

applied to the cases under discussion, they tilt the balance in the direction of the consistent circumstance-involving description. Perhaps needless to say, this invocation of a *ceteris paribus* preference for uniformity allows that inquiry might one day reveal salient differences that make Brown's non-uniform option preferable. But until such relevant disanalogies are revealed, it's hard to see why we should prefer to resolve the standoff in a way that needlessly and positively flouts the standard norm of uniformity (hence, therefore, normal canons of intellectual inquiry).¹⁵

References

- Brown, Derek (2015), "Colour layering and colour relationalism," *Minds and Machines*, vol. 25, pp. 177-191.
- Cohen, Jonathan (2004), "Color Properties and Color Ascriptions: A Relationalist Manifesto," *The Philosophical Review*, vol. 113, 4, pp. 451-506.
- (2008), "Color Constancy as Counterfactual," *Australasian Journal of Philosophy*, vol. 86, 1, pp. 61-92.
- (2009), *The Red and The Real: An Essay on Color Ontology*, Oxford University Press, Oxford.
- Friedman, Michael (1974), "Explanation and Scientific Understanding," *The Journal of Philosophy*, vol. 71, 1, pp. 5-19.
- Kitcher, Philip (1989), "Explanatory Unification and the Causal Structure of the World," in *Scientific Explanation*, ed. by Philip Kitcher and Wesley Salmon, Minneapolis: University of Minnesota Press, vol. 8, pp. 410-505.
- Tye, Michael (2012), "Cohen on Color Relationalism," *Analytic Philosophy*, vol. 53, 3, pp. 297-305.

¹⁵An earlier version of this paper appeared as a response to Brown's contribution to the 2015 Minds Online Conference (<http://mindsonline.philosophyofbrains.com/>). Many thanks to Derek Brown, Matthew Fulkeron, and an anonymous referee for helpful discussion and comments that have much improved the paper.