Projectivism Without Error[1]

Andy Egan, Rutgers University
September 24, 2009

1. The Projectivist Thought

Here is a potentially attractive thought to have about perception:

> Part of the content of our perceptual experience comes from the world, and part of it comes from us. Some of the features attributed in perception are *discovered* – they are objective, observer-independent properties of the objects in our environment, and the reason we attribute them to the things that we do is that our perceptual mechanisms make us appropriately sensitive to the objective features of objects out in the world. Some of the features we attribute to objects in perception are *projected* – their attribution is a product, not of our perceptual openness to the objective features of things as they are in themselves, but of the peculiarities of our own particular sensory apparatus. We, or our perceptual systems, don't just take in the features of things in our perceptual neighborhoods – we also "gild and stain"[2] them with perceptually-attributed features whose origins in perceptual experience have, in fact, more to do with us than they have to do with the objects they're attributed to.

[1] Thanks to audiences at the Australian National University and the SPAWN, and in particular to Bence Nanay, Susanna Siegel, Brad Thompson, Christopher Hill, Frances Egan, and Jonathan Cohen, for very helpful comments and questions.

[2] In Hume's 1751/1983 (appendix I, part V) phrase.

This is the thought that I'll be hoping primarily to cash out, and secondarily to justify, in what follows. It should be a familiar sort of thought. I take it to be a (the?) central *projectivist* thought – the thought that some of the features we attribute to things aren't "really out there" to be discovered, but are instead projected out onto the world by us, or by our perceptual systems. It's a thought that I find quite attractive, but also quite puzzling.[3]

A lot of different views have sailed under the flag of "projectivism" over the years. I'm not going to undertake the impossible task of satisfying every motivation anyone has ever had for offering such a view. All that I will be after here is an account that allows us to cash out the attractive thought above, in a way that makes sense, and that has some chance of being true.

To help give a better sense of the target idea, here are a few places in philosophical history where one can find something like the projectivist thought:

> John Locke:
>
> > The particular bulk, number, and motion of the parts of fire, or snow, are really in them, whether anyone's senses perceive them or no: and therefore they may be called real qualities, because they really exist in those bodies. But light, heat, whiteness, or coldness, are no more really in them, than sickness or pain is in manna. Take away the sensation of them; let not the eyes see light, or colors, nor the ears hear sounds; let the palate not taste, nor the nose smell, and all colors, tastes, odors, and sounds… vanish and cease…
> >
> > (Locke 1690: Book 2, chapter 8, section 14.))

[3] The projectivist thought, as characterized above, has a lot in common with the thought that I think lies behind (a lot of the appeal of) the distinction between primary and secondary qualities. It won't be surprising, then, that what I say about projected qualities below is pretty much what I say about secondary qualities elsewhere (Egan 2006a).

David Hume:

…'tis a common observation, that the mind has a great propensity to spread itself on external objects…

(Hume 1740/1978, I.iii.XIV)

…taste has a productive faculty, and gilding and staining all natural objects with the colours, borrowed from internal sentiment, raises in a manner a new creation…

(Hume 1751/1983: 88.)

William James:

We conceive a given reality in this way or that, to suit our purpose, and the reality passively submits to the conception… In all these cases we humanly make an *addition* to some sensible reality, and that reality tolerates the addition.

(James 1907: 251-252.)

Gideon Rosen:

Successful thought amounts to the *detection* of something real, as opposed to a *projection* onto the real of our own peculiar or subjective perspective.

(Rosen 1994: 278.)[4]

I'll have two goals in this paper. Primarily, I'm interested in laying out a view that's got a good claim to cashing out the projectivist thought, and that avoids some of the costs of other

[4] Rosen isn't speaking with his own voice here – he's articulating what he takes to be an attractive sort of thought that can't, at the end of the day, be given a satisfactory theoretical underpinning. One of the purposes of this paper, and of Egan 2006a, is to provide the sort of underpinning that Rosen's paper argues that we can't have.

ways of cashing it out. (In part, because it avoids the commitment to systematic error that seems to go along with projectivism.) Secondarily, I'm interested in arguing that the proposed version of projectivism is, plausibly, the *correct* account of some of the projectivism-motivating perceptual phenomena. The first project is, I think, still interesting even if the second doesn't succeed – it's helpful to have the best versions of even the views that we don't ultimately want to endorse out on the table.

I'll first survey some of the phenomena that might motivate the projectivist thought, and then look at some ways of cashing out just what it would amount to for the thought to be correct. I'll worry a bit about some of the standard ways, and then advocate another way of cashing out the thought that I think, at least in some cases, does a better job of both capturing the phenomena and underwriting the projectivist idea.

2. A Motivation for Projectivism

One way to motivate the projectivist thought is to look at cases where we find what looks like interpersonal variability in the content of perception – in which different people seem to perceptually represent the same object in incompatible ways – but where we don't want to attribute asymmetric error. (Because, for example, we don't have any good, non-question-begging grounds for saying one party's in a better position to track objective facts in the relevant domain than the other.) One attractive response to this sort of situation is to go projectivist, and say that, in fact, *neither* party is really tracking the facts about the world – both parties are, instead, gilding and/or staining the world (in incompatible ways) with features that are the products of their own sensory particularities.

There are a number of places where there's at least a *prima facie* appearance of being in this sort of situation, where we've got perceivers systematically attributing incompatible-looking features to things in their environments, and no principled basis for attributing error to one party rather than the other. For example:

(i) Gustatory qualities such as *sweetness* and *bitterness* look like good candidates for projectivist treatment, given the variation we see across individuals in how things taste to them, and the lack of a plausible way of picking out a privileged bunch of tasters who are *really* getting it right.

(ii) Aesthetic properties, too, seem – or anyway, have seemed to many – to be good candidates. One often hears people say, for example, that "beauty is in the eye of the beholder", or that "…no sentiment represents what is really in the object",[5] and "to seek in the real beauty, or real deformity, is as fruitless an enquiry, as to pretend to ascertain the real sweet or real bitter".[6]

(iii) Visual-field shapes and sizes – for example, the elliptical appearance presented by the circular (and, indeed, still also circular-looking) plate when it's held at an angle to the eye – are also plausible candidates. There is (one wants to be able to say) some aspect of the way my visual experience represents the plate when I'm looking at it edge-on that's incompatible with an aspect of the way your visual experience represents it when you're looking at it face-on. But it's not attractive to say that, because of this incompatibility in how we represent the

[5] Hume 1757/1965, section 8.
[6] Hume 1757/1965, section 8. I say a lot more about the aesthetic case in Egan (forthcoming)

plate, it's likely that one of us is representing it correctly while the other is misrepresenting it.[7]

(iv) Perhaps also the sorts of affordances one finds in certain bits of psychological theory, and (perhaps) also in philosophers such as Heidegger[8] – when objects are presented to us as *to be used* in certain ways, or as offering certain possibilities for how we might engage with them; and similarly, appearances of importance and salience, when objects or features of the environment present themselves as particularly important, relevant, or interesting. In many cases where some object offers different affordances to me than it offers to you, it's not very attractive to say that one of us is tracking the *real* affordances of the object in question, while the other is getting it wrong.

This is a diverse list, and obviously there will be important respects in which the types of qualities on it differ from one another. What I want to draw attention to, though, is what they have in common: All of these look like candidate cases of perceptual incompatibilities without a plausible basis for an asymmetric attribution of error. (Note also that while I'm focusing on interpersonal variation here, intrapersonal variation will do just as well. Cases of a single perceiver attributing incompatible features to the same thing at different times, in which we don't

[7] There's a lot of discussion of this phenomenon of, to coin a phrase, shape-constancy-while-still-looking-different-in-a-shapey-sort-of-way, and the related phenomenon of color-constancy-while-still-looking-different-in-a-colory-sort-of-way. In this volume, Mohan Matthen's and Sean Kelly's papers take up these issues in some detail. The quick motivational remarks above depend on finding attractive an account of shape-constancy phenomena that's like the account of color-constancy favored by Sydney Shoemaker (2006), according to which (for example) my visual experience always represents the plate as circular, but also attributes some other features to it which explain why it looks different when it's held at different angles. These other features are the ones that I think are good candidates for projectivist treatment.

[8] See for example Heidegger 1927/1962, Gibson 1977, 1979, and Noë 2004.

want either to attribute asymmetric error, or to say that the object underwent a relevant change in the intervening time, are also likely to make a projectivist story look attractive.)

I propose to focus on another example: perceptual representation of *unique*, *pure* or *true* colors.[9] Asked to pick out which of a series of color chips are, for example, *unique blue* – blue with no tinge of any other color – different subjects reliably pick out different ones. Let's simplify and pretend that the only interpersonal differences in unique-hue perception that there are are gender-based. (Of course, the actual pattern of interpersonal variation is more complicated. See in particular Hardin 1988 and Block 1999.) Women, let's suppose, reliably pick chip number 17 of some series as unique blue, while men reliably judge chip 17 to be greenish-blue, and select chip 15 as the unique blue one.

Here we have a case that fits our specifications – we have two groups of people making (apparently) incompatible perceptual judgments, and it doesn't seem plausible to say that there's some bunch of objective facts that they're both making judgments about, and one group is just doing a better job of tracking those objective facts than the other.[10]

It's plausible to think that, when our perceptual systems are dealing in attributions of unique hues, this isn't a matter of sensitivity to the genuine, out-there-to-be-discovered property of e.g. *unique blueness*, which some of us are correctly tracking and others not. Instead (the thought continues), that bit of perceptual content is due to idiosyncratic features of our own particular representational apparatus. Stuff that's really about us and our perceptual equipment is

[9] See Hardin 1988, Block 1999, Tye 2006a, 2006b, 2007, Byrne and Hilbert 2007, Cohen, Hardin and McLaughlin 2006 for extensive discussion.

[10] I'm not going to be very careful about the distinction between the contents of perceptual *experiences* and the contents of perceptual *judgments*. I'll write as if it was uncontroversial that the contents of our perceptual judgments (almost) always just reflect the contents of our perceptual experiences. That may not be right, but even if it's not, the added complications won't, I think, make a difference to the argument here.

getting projected out onto the world. What attributions of qualities such as unique blueness are really responsive to, it's attractive to say, is features of the individuals doing the representing and their perceptual systems, and not genuine, objective features of the objects being perceptually represented.

3. Cashing out the Metaphor

These sorts of cases are ones in which it's appealing to think that something like the projectivist thought is correct. Which raises the question: Just what would it amount to for the projectivist thought to be correct, exactly? What would have to be going on in perception for that sort of "gilding and staining" talk to be appropriate?

Before we go on, let me note some ground rules that I'm going to impose on the kinds of stories that will be under consideration:

I'm going to look at ways of cashing out the projectivist idea as a distinction between the different sorts of qualities or features that we attribute to things in perception – so what I'll be after is a story about what some perceptually-attributed feature has to be like in order for the projectivist thought to be correct about it.

The story about what the relevant features are like is going to be told in terms of their contributions to the veridicality conditions of experiences – in terms of what it would take for an attribution of such a feature to be correct. On the sort of picture I'll be working within, the attribution of a feature to some entity determines a veridicality condition. Indeed, the main business, for our purposes, of a feature (such as *being a dog*) is to determine a veridicality condition when it's attributed to an object (such as Lassie) in an experience (or a thought, or

whatever). So, I'll be after a distinction between kinds of features attributed in experience, in terms of the veridicality conditions of their attributions, that's fit to underwrite the projectivist thought, by providing a satisfying distinction between the projected features and the rest.

I'm also going to understand veridicality conditions in possibility-carving terms – in terms of the way that they divide up a class of possibilities into the ones in which things are as represented in the experience, and the ones in which things are otherwise.

Since we'll also (I hereby stipulate) be thinking about the objects of thought in possibility-carving terms, determining a veridicality condition is going to be sufficient for determining an object of thought. So what an attribution of a feature to an object does is determine an object of thought – a candidate thing to believe, desire, etc.

Working in this sort of possibility-carving framework will allow me to move smoothly back and forth between talking about the *correctness conditions* of perceptual experiences and talking about the *contents* of perceptual experiences. (There are, of course, notions of content that wouldn't allow such smooth movement from one to the other, but we won't be employing them here.)

We can think of the features that get attributed in perception, and in thought, then, as functions from entities to potential objects of thought – from the entities that they might be attributed to to the objects of thought that would be determined by such an attribution. This should be a pretty familiar sort of picture.

All of these restrictions are partly in the service of keeping the discussion manageable, and partly in service of restricting the theoretical commitments of the resulting projectivist account to ones that I'm comfortable with. I think that we're going to wind up committed to

veridicality conditions for experiences come what may.[11] But I'd prefer that my account of the projectivist thought not force me to take sides in disputes about the details of the representational mechanisms by means of which those veridicality conditions get determined.

By imposing these constraints, I am taking a number of candidate accounts off the table – I'm ruling out, for example, mode-of-presentation based views that put important theoretical weight on non-possibility-carving aspects of representation, and I'm ruling out the sort of expressivist projectivism advocated by Simon Blackburn (in, for example, Blackburn 1984) which doesn't trade in possibility-carving at all. I'm doing this not because I take myself to have some great in-principle objection to such accounts, but because I want to explore the prospects for a view that satisfies the constraints just set out. I've got theoretical and methodological proclivities that make me want to stick to stories told in these kinds possibility-carving terms until I'm forced out of them, but even if you don't share those proclivities, there's value in exploring how much we can do with these resources.

I'll conduct the discussion in terms of the unique hues, and cast the task of satisfactorily cashing out the projectivist metaphor as the task of providing a reasonably appealing projectivist account of the perceptual phenomena surrounding our judgments about e.g. unique blue. The hope is that the account of the unique hues will serve as a good illustration of the general projectivist strategy I'm advocating, which could then be implemented in other cases. (And indeed, could still be implemented in other cases even if we decide, at the end of the day, that it's not the right story to tell about the unique hues.)

[11] Not everybody agrees. But see Siegel (this volume) for some support.

It will also make presentation simpler to have a particular case to focus on. How to generalize the account of the case should be clear enough.

Let's suppose that Ron and Hermione look at a series of color chips. Hermione's visual experience represents chip 17 as unique blue. Ron's visual experience represents chip 17 as greenish-blue – as blue slightly tinged with green – and represents chip 15 as unique blue.

Suppose that we want to give a projectivist account of this bit of the content of Ron's and Hermione's visual experience, in order to account for their perceptual variation without saying that one, but not the other, is correctly tracking the unique hues of the objects in their environment. What should we say, exactly, about the ways in which their respective experiences represent the chips? In particular, what should we say about the qualities *unique blue* and *greenish blue* that Hermione's and Ron's experiences respectively attribute to chip 17?

Let's start with some desiderata that we'll want our projectivist account of the unique hues to satisfy:

INCOMPATIBILITY: The features attributed in Ron's and Hermione's experiences are genuinely incompatible. One couldn't have a veridical experience that attributed both to the same object.[12]

NO ASYMMETRIC ERROR: It's not the case that one of Ron or Hermione is getting it right, while the other is getting it wrong.

[12] There are a lot of ways to cash out the notion of incompatibility. I don't claim that the incompatibility requirement that I focus on here – about the impossibility of a single perciever's veridically representing the thing both ways simultaneously – is the only respectable notion of incompatibility. It's not. What's important to making the argument here go is just that it's *a* legitimate notion of incompatibility, such that satisfying the INCOMPATIBILITY requirement is a way of underwriting the intuitive appearance of incompatibility.

PROJECTION: What's going on warrants talk of projection – of gilding and staining, of being the contribution of us and our perceptual apparatus rather than (just) the way the world is, etc. The qualities attributed (*unique blue* and *greenish-blue*) have got to be such that their attribution is traceable to features of the subject's particular sensory apparatus in a way that attribution of other sorts of qualities isn't.

The intuitive case for INCOMPATIBILITY is, I think, pretty clear and compelling. The way Hermione's visual experience represents chip 17 is just incompatible with the way Ron's visual experience represents it – one couldn't consistently represent chip 17 both as Hermione's experience represents it and as Ron's experience represents it. So, being *unique blue* had better preclude being *greenish-blue*, and v.v.. (More carefully: Hermione's correctly representing chip 17 as unique blue had better preclude her simultaneously correctly representing it as greenish-blue.)

We want NO ASYMMETRIC ERROR because neither Ron nor Hermione seems to have a better claim than the other to be tracking the genuine facts of the matter. There doesn't seem to be any relevant difference in the way their perceptual systems are responding to the things in the world that they represent that would warrant the claim that one, but not the other, is having a veridical experience of the color of chip 17.[13]

[13] Byrne and Hilbert (2007) offer an account that embraces asymmetric systematic error. On their view, somebody's getting it right, but we're just in a lousy position to say who, since we aren't privy to the causal and evolutionary facts that fix the content of the relevant visual experience types. There's a perfectly good *metaphysical*

PROJECTION promises to explain how their perceptual experiences come by a sort of content that isn't tracking genuine, objective features of objects in the world – it's there on account of the features of the perceivers and their perceptual apparatus, not on account of objective features of the things perceived.

Note that I'm not claiming here that these desiderata are non-negotiable constraints on the theory of unique-hue perception that we'll want, at the end of the day, to endorse. I don't think that they are. What I do think – and what's important for our purposes now – is just that they're (at least partially) constitutive of a sort of view that has some appeal, and that it would be nice to be able to have a plausible version of on the table.

We'll look first at a few standard-ish options for how to flesh out this sort of view. I'll also complain about them a bit in order to motivate the rival account that I'll offer in the next section. I'm not aiming for an argument by elimination here – I certainly won't be looking at all of the available options, and I won't be aiming to conclusively rule out the ones that I do look at. Rather, I'll be using some selected rival theories as points of contrast in order to highlight what I take to be attractive features of the sort of view that I'm going to offer.

One option is what Sydney Shoemaker (1990) calls *literal projectivism*: There are some properties that are in fact just properties of our experiences, but which we, and/or our perceptual experiences, mistakenly attribute to things outside the mind. On this kind of view, *unique blue* and *greenish-blue* are actually qualities of Ron's and Hermione's perceptual experiences, which they, or their visual perceptual systems, mistakenly attribute to objects out in the world.

---

basis for the asymmetric error, but we're in a lousy position to adjudicate the dispute. I don't have any very good arguments that this *couldn't* be the right story, but I think it's an unattractive thing to say, and its unattractiveness ought to incline us to look for other options.

Another is what Shoemaker calls *figurative projectivism*: There are some additional properties, instantiated neither by our experiences nor the objects in the world that we attribute them to, that we attribute to things on account of having certain sorts of experiences. On this sort of view, *unique blue* and *greenish blue* aren't features of Ron's and Hermione's visual experiences, but neither are they features that are actually possessed by the objects to which they are attributed.

Assuming that we fill in the details right, both of these sorts of accounts will satisfy all of the desiderata above. It's easy to specify the properties attributed in such a way that they're incompatible, and so that nothing (or nothing outside the mind) has the properties either party's perceptual experiences attributes to the color chips. So both parties' perceptual experiences are nonveridical, and there's no asymmetry in our attributions of error. And these are pretty clearly projectivist views – either features of our perceptual experience are themselves being attributed to objects in the world, or else other sorts of properties are getting attributed to objects in the world, not on account of our perceptual faculties being receptive to the presence of those properties in the objects we perceive, but (merely) on account of some of the properties of our perceptual experiences of those objects.

Still, these accounts still aren't completely satisfying. Both of these options convict visual perception of a particularly bad sort of systematic error – either confusing properties of experience and properties of things out in the world, or else conjuring properties out of whole cloth to attribute to things out in the world. These are both pretty serious and surprising sorts of error.

It shouldn't be surprising that these are both error theories, since given our desiderata, an error theory actually looks kind of inevitable. When we're faced with two perceptual judgments, it looks as if there are only three possible things to say about their veridicality: (1) they're both veridical, (2) they're both non-veridical, or (3) one is veridical and the other is non-veridical. The ban on asymmetric error rules out (3), and the incompatibility requirement rules out (1), which leaves (2) as the only option on the table. So it's likely to seem as if, given INCOMPATIBILITY and NO ASYMETRIC ERROR, we're just going to be stuck with an error theory. Still, it would be better if this could be avoided. (I'll argue in the next section that it can be.)

There are several reasons to be reluctant about signing up for an error theory. One is just the usual sort of application of some (fairly weak) charity principle as a constraint on interpretation. Other things equal, it's better to avoid attributing systematic error to the subjects and systems that you're interpreting.

Another motivation is the reason-givingness of the relevant bits of perceptual content. It seems as if, for example, gustatory perception of objects as sweet, salty, bitter, etc. gives us good reasons to eat one sort of thing rather than another. Having a visual experience that represents the streetlight as red gives me a good reason to stop the car. And so on. If those experiences are never veridical, since they're trading in properties not possessed by things outside the mind, we'll need to do some work to say that they're nonetheless appropriately reason-giving, and our account of their reason-givingness is liable to be less direct, and more convoluted, than it seems like it ought to be.

Finally, there are some veridicality-ish distinctions in the neighborhood that we'll definitely want to be able to make, but which a global error theory threatens to mask. Even if we

don't want to make a distinction in terms of veridicality between my attributions of unique-blueness in good viewing conditions and your attributions in good viewing conditions, we *do* want to make such a distinction between, for example, my attributions in good perceptual conditions and bad, my attributions where I'm just guessing with my eyes closed and (we'd like to say) I get it wrong, and the ones where I'm looking in good lighting conditions and (we'd like to say) I get it right.

Compare the case in which Ron, eyes closed, makes a wild guess about chip 2 (which, if he was looking carefully in good conditions, he'd say was unique red), and judges that it's unique blue, to the case in which he looks, in good conditions, at chip 2 and judges that it's unique red. We probably want to say something positive, in a veridicality-ish neighborhood, about the second judgment that we don't want to say about the first. Similarly in the case where Ron looks at (what we'd normally be inclined to say is) a white chip, under an appropriately selected sort of blue light, and has a visual experience that represents it as unique blue, as compared to the case in which he looks, in paradigmatically good viewing conditions, at chip 15, and has an experience that represents it as unique blue. Another example of the same sort of phenomenon: Asked to put all of the unique-blue things in a pile, there are ways of doing better and worse at the task. We want to be able to draw a distinction between the case in which Hermione winds up with a pile that includes chip 15 and a bunch of appropriately similar blueberries, and the case in which she winds up with a pile that includes, say, chip 17, chip 2, and a bunch of bananas. A general error theory, according to which nothing at all is, for example, unique blue, threatens to mask these distinctions by just consigning *all* attributions of projected qualities to the "non-veridical" bin.

If we could manage it, then, it would best to be able to accommodate a fourth desideratum:

CORRECTNESS: Both Ron's and Hermione's perceptual attributions of *unique blue* and *greenish-blue* are correct.

We could accommodate this desideratum by adopting another popular sort of view – some variety of response-dependent contextualism or externalism, according to which they're each just tracking different objective properties of things under the description or mode of presentation *unique blue*. On this kind of view, there's no such thing as *unique blue* simpliciter. There's just a family of properties, *unique-blue*$_K$ for various kinds K. Males' judgments about what's unique blue are about what's *unique-blue*$_{\text{Male}}$, while females' judgements about what's unique blue are about what's *unique-blue*$_{\text{Female}}$, and so on for whatever the relevant groups turn out to be. (There are different ways to spell out the details of such a view. One, dispositionalist, way is to say that x is unique-blue$_K$ iff x is disposed to cause sensations of a certain phenomenally-individuated type – call it B – in Ks in normal viewing conditions. Another, physicalist, way to cash it out is to say that unique-blue$_K$ is that physical property, whatever it is, that forms the categorical basis of the disposition to cause B sensations in Ks in normal viewing conditions. This is, obviously, just the beginning of a very long list of options.)[14]

[14] I mean to be casting quite a wide net in describing this category of view - there are a *lot* of people who, as I read them, are offering views of this general type, though not all of them describe their views in quite these terms. Such views exhibit a lot of variety in the sorts of properties that they take to be represented by the relevant types of experiences, and in the sorts of features of the perceiver and/or her perceptual environment that they take to be relevant to determining just which of the selected class of properties a given experience-type is going to represent. But these differences, while vitally important for very many purposes, don't matter for ours.

A nice feature of this is that it allows us to say that everybody's getting it right, because everybody's perceptual experiences are dealing in different bunches of properties, and everybody's correctly (more or less) tracking their own proprietary bunch of properties. In particular, Ron and Hermione are both correctly representing chip 17. Ron's experience correctly represents chip 17 as *greenish-blue*$_{\text{Ron}}$ (or *greenish-blue*$_{\text{Male}}$), and Hermione's experience correctly represents it as *unique blue*$_{\text{Hermione}}$ (or *unique blue*$_{\text{Female}}$).

Putting it in terms of concepts deployed in perception, we can say: There's a perceptual concept, unique blue, which different perceivers are deploying differently, but since they each deploy it to track different properties, everybody's deploying it (more or less) correctly.

If we don't like conceptual content in perception, we can still say: There's a common way they're both deploying their representational apparatus, but in response to different stuff – there's a representationally significant feature of the insides of their heads that they're both deploying, but in different ways. There's some visual representational state R such that Hermione gets into R in response to chip 17, and Ron gets into R in response to chip 15. But

---

The color pluralist views defended in, for example Kalderon (2007), Allen (2005), Bradley and Tye (2001) and Mizrahi (2006) are pretty clearly of this type. As I read Brian McLaughlin (2003), for example, he's also offering a view of this type, according to which my unique-green-ish-experiences in circumstances C1 and your unique-green-ish-experiences in circumstances C2 are (very likely to be) attributing different physical properties to the objects in our environments. Another sort of view of this type is the sort of relationalist view (or family of views) defended by Jonathan Cohen (in, for example, Cohen 2004, 2006, 2007, forthcoming), on which humans' and pigeons' visual experiences, and the visual experiences of different humans with appropriately different visual systems, are in the business of representing different classes of relational properties. Color dispositionalist views (such as the ones we find in Jackson and Pargetter (1987) and Mark Johnston (1992)), which hold that members of different perceptual kinds are (likely to be) attributing different dispositional properties to objects in their phenomenologically similar visual experiences will also fall into this category. So too will the views of philosophers who, like Brad Thompson (2007, forthcoming), advocate a Fregean view of perceptual content according to which a common phenomenal mode of presentation can be deployed, in the experiences of different sorts of perceivers, to represent different properties of objects in their environments.

since that representational state is tracking different properties in each of them, they're both deploying it/instantiating it/whatever (more or less) correctly.

I'll complain about views of this sort below, but the complaints I make won't, primarily, be arguments that they're not *true*, so much as arguments that they're not good ways to capture the projectivist thought. Of course it could turn out that we ought, at the end of the day, to reject the projectivist thought. But whether we want to endorse a projectivist view at the end of the day or not, it would be helpful to know just what a projectivist account would look like.

There are two reasons to be concerned about whether this kind of view really delivers on the projectivist thought:

First worry: It doesn't give us any incompatibility in the *content* of Ron's and Hermione's perceptual experiences. There's no incompatibility because there's no common subject matter – they're not really making conflicting or incompatible judgments, since their perceptual mechanisms are just dealing in attributions of different properties.

Response: It's true that we don't get incompatible contents – that is, we don't get incompatible correctness conditions. We do, however, get a different sort of incompatibility – we get incompatible ways of representing at the level of something like modes of presentation, or perceptual states. There's a perceptual concept, or a way of perceptually representing, that one applies and the other withholds. (Perhaps better, there are incompatible perceptual concepts, or incompatible ways of perceptually representing, that they apply.)

It's not clear that this is incompatibility enough. It would be better, I think, to be able to deliver the story at the level of correctness conditions if we can. First, because it seems to do

fuller justice to the intuition of incompatibility. And second, because it doesn't run the risk of embroiling us in a big fight about the mechanisms of perceptual representation.

Absent fancy footwork (which is coming in the next section), CORRECTNESS looks to be incompatible with INCOMPATIBILITY – at least, with incompatibility at the level of veridicality conditions. And so we'll have to choose one of the initially appealing desiderata to give up. Shoemaker's literal and figurative projectivisms preserve INCOMPATIBILITY while giving up CORRECTNESS. As a result, they buy themselves some theoretical work in explaining the reason-givingness of perception of unique hues, and the apparent distinctions in terms of veridicality to be drawn between different unique-hue perceptions. The sort of response-dependent contextualist view just considered preserves CORRECTNESS while giving up INCOMPATIBILITY. As a result, it buys itself some theoretical work in explaining the appearance of incompatibility.

Whatever we think about the not-enough-incompatibility complaint, there's a second worry which is, I think, more important for our purposes. Whatever there is to be said for this sort of contextualist/externalist story as an end-of-the-day theory of the phenomena, it doesn't really look good as a way of cashing out the distinctively projectivist idea. This is a theory on which that what's going on is that there's one bunch of perfectly objective properties that group A is tracking, and another bunch of perfectly objective properties that group B is tracking – in our case, Ron's attributions of unique blueness are tracking *unique blue*$_{\text{Ron}}$ (or *unique blue*$_{\text{Male}}$), while Hermione's are tracking *unique blue*$_{\text{Hermione}}$ (or *unique blue*$_{\text{Female}}$). This sort of view is probably better described as *filterism* than *projectivism*. All the properties that feature in perception are genuine, objective properties to be found "out there", which we (typically) attribute to things on account of their genuine, objective presence in the things to which we

attribute them. It's just that which of the relevant properties our particular sensory faculties are in the business of representing is a contingent, and potentially idiosyncratic, fact about us.

This doesn't look like a good way of capturing the thought that the qualities in question are contributed by us, rather than being "out there" to be discovered. On this sort of view, the 'projected' bits of perceptual content are still instances of content that's there as a result of openness to objective facts about the world – it's just that different ones of us are perceptually open to different bunches of objective facts.

In addition to not really seeming to warrant talk of *projection*, this sort of account isn't going to provide an interesting distinction between different bits of perceptual content, which allows us to distinguish the projected bits from the rest. For *any* feature of objects in our environment that we're perceptually sensitive to, it's a contingent fact about us that we've got perceptual systems that are sensitive to that feature. So this fact about some perceptually represented quality – that it's on account of some contingent peculiarities of our particular sensory apparatus that our perceptual systems are in the business of tracking that quality rather than some other – doesn't separate the things about which projectivism is appealing from the things about which it's not in the right sort of way. Since *everything* counts as projected by this standard, we don't get the right sort of distinction between the features that are supposed to be projected and the ones that aren't.

4. Another Option

I'm going to advocate a view according to which perceptual experiences typically have self-locating (also known as *de se*) contents – that is, self-locating veridicality or correctness

conditions. At a first pass, the features that warrant projectivist treatment are those whose attribution to objects is responsible for the self-locatingness of perceptual content. (A slightly fancier version of the view will be coming later.) Some explanation of this is in order.

Perceptual experiences have *veridicality conditions*. Recall that we're thinking of these as the classes of possibilities in which things are as the perceptual experience represents them to be. The veridicality conditions of an experience are determined by how it attributes features to objects in the perceiver's environment. Attributing some feature to an object determines a veridicality condition – a class of possibilities. A perceptual experience that attributes *squareness* to chip 17, for example, will be veridical only in possibilities in which chip 17 is square. Perceptually attributed features are the sorts of things that determine functions from objects to these sorts of possibility-carving veridicality conditions.

Let's simplify our picture of perceptual content a bit, and suppose that all perceptual content consists of attributions of features to particular objects in the perceived environment. Each attribution of a feature to an object determines a correctness condition. A perceptual experience as a whole is veridical in just those possibilities where all of the objects represented have all of the qualities attributed to them. (This is the intersection of the classes of possibilities determined by each of the feature-attributions that makes up the total content of the experience.)[15]

[15] I'm supposing, for presentational convenience, (a) that visual perception has object-involving content – content that attributes features to particular objects in the environment, rather than just representing, for example, that it there is some object with the represented feature (contents it's natural to think of has being of the form *Fa* rather than $\exists x(Fx)$), and (b) that visual perception has *only* object-involving content. Now, (a) is contentious, and (b) is certainly false. I'm making these assumptions because they make for a simpler picture of the content of perception, on which it's easier to talk about how the attribution of features to objects in the environment determines correctness conditions. But we could tell a very similar, though slightly messier, story of the way that the features attributed in perception contribute to, for example, existential rather than object-involving correctness conditions.

What kinds of things are these *possibilities*? An initially attractive, and fairly standard, thing to say is that they're possible ways for the world to be, or *possible worlds*. The class of possibilities picked out by Hermione's perceptual attribution of *squareness* to chip 17 is the class of worlds in which chip 17 is square.

There is, however, good reason to think that the possibilities that we distinguish between in *thought*, at least, are finer-grained than this. When we have beliefs about what time it is, about who we are, or about features of our own particular predicament, for example, we take a stand on more than just the global facts about the world is like – on more than just which world is actual.[16]

When I believe that it's noon, the accuracy of my belief doesn't just depend on which world is actual – it also depends on which time is present. Since fixing which world is actual doesn't fix which time is present, the way in which my beliefs about the time distinguish between possibilities is not well-captured by thinking of their contents as sets of possible worlds. When I have a belief about the time, I take a stand not just on which world I inhabit, but (also) on my temporal location within it. My beliefs about the time distinguish not between *worlds* but between *locations within* worlds.

Other cases show us that this phenomenon isn't specific to times. Here is a variation on a familiar sort of example: Harry Potter is lost in Hogwarts School of Witchcraft and Wizardry after an amnesia-inducing magical accident. It could be that he's lost because he's missing some information about what the world is like – that is, because he doesn't know which world is actual. He could be lost, for example, because he doesn't know what the actual floor plan of the

[16] The discussion of self-locating content that follows is my summary of the upshot of an extensive literature, including Castañeda (1966, 1967, 1968), Lewis (1979), Chisholm (1982), Sosa (1983), Peacocke (1992), Noe (2004), and others. Perry (1979) famously dissents from the conclusions I (along with e.g. Lewis and Chisholm) want to draw from the phenomena that I take to motivate a self-locating picture of representational content, as does Stalnaker (1990).

castle is. But Harry could still be lost, even if he knew everything there is to know about which world is actual. Let Harry be looking at the Marauder's Map, which shows the complete floor plan of the castle, as well as the location of all of its inhabitants. He doesn't, then, lack any relevant information about what the world is like. He could, nonetheless, still fail to know where *he* is, because he could still fail to know which, of all of the creatures in the castle, is *him* – perhaps the accident robbed both Ron and Harry of their memories, and now Harry is unsure whether he is amnesiac-Harry or amnesiac-Ron. Harry knows all of the relevant facts about the world – he knows, for example, that Harry Potter is in the West wing, and that Ron Weasley is in the East wing – but he still doesn't know whether *he* is in the West or the East wing, because he doesn't know whether he is Harry or Ron.

What Harry is ignorant about in this case isn't what the world is like – Harry knows all of the relevant facts about the world. What Harry is ignorant about is something about his *location within* a world about which he is already as well-informed as he could hope to be. In order to remedy his ignorance, Harry needs to rule out some possibilities. But the possibilities he needs to rule out aren't possible ways for the world to be – they're possible locations, situations, or predicaments within a world. (In this case, the possible predicament that Harry needs to rule out, but hasn't yet, is Ron's.)

So, to model all of the kinds of belief and ignorance that we want to model, we need our possibility space to be a space not of possible *worlds*, but of possible *predicaments*.

If the contents of *belief* ought to be modeled in terms of a space of possible predicaments rather than a space of possible worlds, then so ought the contents of perception. For one thing, we want our theories of belief and perception to play nicely with each other. We want, for

example, to be able to say that people often come to believe that things are as their perceptual experiences represent them to be – that is, that people often come to have beliefs with the same contents (i.e., the same correctness conditions) as their perceptual experiences. (At least: sometimes we come to have a belief that P because we have a perceptual experience that represents it as being the case that P. Maybe we never come to believe the *whole* content of any perceptual experience, perhaps because perceptual experiences are too informationally rich for all of their content to make its way into belief.)

Second, it should already have been clear that perception doesn't just deal in the sort of God's-eye-view information about which world is actual that's happily modeled in a possible-worlds framework. Visual perception, for example, tells us things about our location, not just God's-eye-view stuff about how objects are arranged relative to one another.[17] That's a good thing – just knowing the possible-worlds-y stuff isn't enough to guide action. For pretty much any world you pick, there will be a lot of different predicaments to occupy within it. The occupants of some of those predicaments would be well advised to duck. The occupants of other predicaments would be well advised to jump. Some of the people who find themselves in hedge mazes would be well advised to turn left, others to turn right. Some of the inhabitants of the actual world are confronted with glasses of gin, while others are confronted with glasses of petrol. So, just knowing which world is actual won't tell you whether it's a good idea to duck, to jump, to turn left, to take a drink, etc. In order to make good decisions about what to do, you

[17] Others who say this, in different terminology, include Peacocke (1992), Schellenberg (2007, 2008), Noe (2004), and Millner and Goodale (2008), to name a few. Talk of *situation content* or *egocentric content* in perception, or of perception as serving to locate us in *egocentric space*, are all instances of this phenomenon.

need to know not just stuff about what the world is like, but also stuff about the nature your own particular predicament.

So: at least some of our thoughts, and at least some of our perceptual experiences, have a sort of content that distinguishes between possible predicaments, not (just) possible worlds. Call such content *self-locating* content, to contrast it with *possible worlds* content, which distinguishes (only) between possible ways for the world to be. Both kinds of content serve to carve out regions of a space of possibilities – the difference is just in the nature of the points in the possibility space that they carve up. (I'll also talk in what follows about *self-locating propositions* and *possible-worlds propositions*, by which I'll mean, respectively, classes of possible predicaments and classes of possible worlds.)

An important cautionary aside: It's important not to get so wrapped up in *spatiotemporal* self-location that we start to think that all self-locating belief is belief about either our temporal or geographical position. When Harry doesn't know whether he's Harry Potter or Ron Weasley, we don't just want to be able to model amnesiac Harry's ignorance about such geographical-location matters as whether he's in the East or West wing of Hogwart's. We also want to be able to model his ignorance about, for example, such non-geographical (and non-temporal) matters as whether he is the child of Arthur and Molly Weasley or the child of James and Lily Potter, whether he has red hair or black, whether he's expected to be in Potions class or Divination class on Tuesday afternoons, etc.

Though a lot of the cases that are used in the literature to motivate a move to self-locating/*de se* content are examples of people who are ignorant of their spatial and/or temporal

location,[18] those sorts of cases aren't the only motivation. The move to a self-locating picture of content can also be motivated by, for example, wanting to find a common object of belief for all of the people who believe that their pants are on fire, or that they themselves are millionaires, and by wanting to find a common object of desire for all of the people who want to be firefighters, president of the United States, or heavyweight champion of the world. None of these sorts of beliefs or desires are happily modeled in terms of a possibility space whose points are possible worlds, and none of them are about spatiotemporal location. So the predicaments in our possibility space are best thought of, not as spatiotemporal locations, but as possible situations for an individual to occupy, which might differ from one another in a host of ways. Some of the important differences and similarities between predicaments will be geographical or temporal, but many will not.

What unifies all of the predicaments that are the focus of Hector's doxastic attention when he thinks the he himself is a millionaire isn't anything geographical or temporal – it's something economic. What unifies all of the predicaments that are the objects of Hillary's and Barack's common desire when they both want to be president isn't anything to do with spatiotemporal location – it's something political. There are, of course, lots of classes of predicaments that are unified by their spatiotemporal features – all of the in-the-oval-office predicaments, and all of the 7:35-pm-EST-on-November-15-2008 predicaments, for example. But in addition to these, there are many other interesting classes of predicaments besides – the presidential predicaments, the millionaire's predicaments, the predicaments of the burning-pantsed – whose members don't have any interesting geographical or temporal commonalities. I

[18] Think, for example, of Sleeping Beauty, Lingens in the library, or Lewis's two gods.

belabor this point because it's very easy to read "self-locating" *much* too narrowly, and to think that self-locating representation must always be all about one's geographical or temporal location. Many of the usual examples used to motivate the idea make this misreading extremely natural. Please don't read it that way. (If you find it too difficult to read "self-locating" in the broader way, it might help if you take me to be systematically misspelling "de se" throughout the paper.)

While we're digressing, a quick technical aside: There are a number of interchangeable ways of talking about self-locating content. I've been talking in terms of sets of possible predicaments (positions, locations, situations…), and treating predicaments as primitive. This is now my preferred way to do things, for reasons that aren't relevant to our purposes here. But it's not quite the standard way of talking. A more common way to cash out the same sort of talk about carving up a finer-grained space of possibilities is to talk in terms of *centered worlds* – triples of a world, a time, and an individual.[19] Abstracting away from some technical metaphysical complications, these two ways of talking are interchangeable. Centered worlds serve the same purpose of picking out a possible situation in logical space for an individual to occupy. (We could, if we subscribe to the right sort of modal and temporal metaphysics, also do the same work using instantaneous time-slices of worldbound possible individuals as the points of our possibility space.) Another way – probably the most common way – to talk about self-locating content is in terms of the self-attribution of properties. (This is the way that David Lewis and Roderick Chisholm talk most of the time in Lewis 1979 and Chisholm 1982.) Given a permissive ontology of properties, these will all be interchangeable, since properties and sets of

[19] These are Lewisian (1979) centered worlds rather than Quinean (1969) centered worlds.

predicaments, or sets of centered worlds, will be in one-one correspondence. So to self-attribute the property, *being in the East wing of Hogwarts* is to believe the self-locating proposition that picks out all and only the *in the East wing of Hogwarts* predicaments – that is, that picks out all and only the world, time individual triples <w,t,i> such that i is in the East wing of Hogwarts at t in w.

And in fact, as Lewis (1979) points out, we can make do with just the one kind of content – for each possible worlds proposition, there is, to introduce a technical term, a *boring* self-locating proposition that, for each world, includes either all or none of the predicaments within it. (Equivalently: there's a boring world-occupancy property, of the type, *being an inhabitant of a world such that*…. Equivalently, there's a centered-worlds proposition that, for every world w, either contains all or none of the <world, time, individual> triples that have w as a member.) So the distinction that's actually doing the work here is the one between *boring* and *interesting* self-locating propositions, rather than the one between possible-worlds and self-locating propositions. But it makes exposition easier and more intuitive to stick with the possible-worlds/self-locating contrast, so I'm going to do that, with the understanding that everything could be rephrased in the obvious way in terms of the boring/interesting distinction. End technical aside.

With this distinction between kinds of contents or correctness conditions in hand, we can make a distinction between kinds of features attributed in belief or perception. We said before that attributing a quality to an object determines a correctness condition – a class of possibilities. Attributing a quality like *squareness* to an object – chip 17, say – will determine a possible-worlds correctness condition – the class of *chip 17 is square* worlds. (Equivalently, we might say it determines a boring self-locating correctness condition – the class of predicaments in *chip*

17 is square worlds.) Lots of qualities will be like this – the correctness conditions determined by their attributions will always be possible-worlds propositions (or boring self-locating propositions).

Once we're working in a framework that allows for self-locating content, we have room for a different sort of quality, whose attribution to objects determines a self-locating correctness condition. (A genuine, interestingly self-locating correctness condition – one that doesn't just divide up predicaments along the boundaries of worlds.) Some plausible candidates include such qualities as *being nearby*, *being on the left*, or *being three feet away*.[20] These are qualities that, applied to an object, determine a set of possible predicaments, rather than a set of possible worlds.[21]

Attributing *being nearby* to chip 17 will determine a class of possible predicaments that is likely to include some, but not all, of the predicaments in a given world – it will include the *near chip 17* predicaments, and exclude the rest. Suppose we both have experiences that attribute *being nearby* to chip 17. Our experiences will share a common correctness condition – the class of *near chip 17* predicaments. If in fact chip 17 is near you and far from me, your predicament will satisfy this shared correctness condition, while mine will not. (Since your predicament, but

[20] Note: by using these English expressions to draw your attention to the relevant qualities, I don't mean to be signing up for the view that, for example, the English word "nearby" standardly has such a feature as its semantic value. (In fact I think it almost certainly doesn't.) In this paper, I'm just concerned with the contents of perception – the issues about language are complicated, and I won't defend any of those claims here. For some discussion of *de se* content in language, see Egan (2007, forthcoming, MS)

[21] Again, there are some other, terminologically but not substantively different, ways that we could put the same point: We could give a uniform account of correctness conditions, and say: all qualities determine sets of predicaments. We'll then say that the centering features are the ones such that the set determined sometimes includes some but not all of the predicaments in a world. We could also characterize attributed features not as functions from objects to sets of points in a possibility space, but as functions from points in a possibility space to extensions. Then we'll say that the centering features are the qualities that determine an extension relative to a predicament, not relative to a world. (Alternatively: that sometimes determine different extensions for different predicaments in the same world.)

not mine, will be a member of the self-locating proposition determined by the attribution of *being nearby* to chip 17.) So your experience will be veridical, and mine not. Still, there will be a common correctness condition for our two experiences, there will be a common feature that both of our experiences attribute to chip 17, and there will be something that we both believe if we both take our experiences at face value. (Though you will be right to believe it, while I won't.)

Let's call qualities of the first sort – the ones whose attribution always determines a possible-worlds proposition, or a boring self-locating proposition – *objective properties*, and qualities of the second sort – the ones whose attribution (at least sometimes) determines an interesting self-locating proposition – *centering features*.[22]

Centering features bear a lot of the distinctive marks of projected features, and self-locating content bears a lot of the distinctive marks of projected content. To illustrate this, let's start with an example of geographical self-location. If Hermione's visual experience represents chip 17 as *square*, it's representing an objective property of the chip, out there to be discovered by any observer who happens to wander by. If her experience represents chip 17 as *three feet away*, it's not representing an objective property of the chip, out there to be discovered by just any old passing observer. It's representing something idiosyncratic about Hermione's particular perceptual predicament. If Harry is standing six feet from chip 17, his visual experience ought, if it's to be veridical, to represent chip 17 as *square*, (and as *three feet from Hermione*) but it ought not represent the chip as *three feet away*.

[22] In Egan (2006a, 2006b) I used *property* and *centering feature* to mark the same distinction.

If Hermione's visual system is working properly, the bit of her visual perceptual content in which chip 17 is represented as *three feet away* will be present in her visual experience not (just) because of her sensitivity to the objective facts about the objects in her environment, but (also) because of the peculiarities of her own particular perceptual situation. If her experience represents the chip both as *square* and as *three feet away*, part of the content of her experience will be dealing in the objective properties of things in her environment, and part will be dealing in something else, having to do as much with the perceiver as it has to do with the objects perceived. Part of the content of her experience – part of the way that she represents chip 17 – will be present on account of Hermione's perceptual systems' openness to the objective facts about the objects represented, and part will be present on account of features peculiar to her.

Now consider both Harry's and Hermione's experiences: Harry represents chip 17 as *six feet away*, while Hermione represents it as *three feet away*. There's a clear sense in which their experiences are *incompatible* – their correctness conditions are, or pick out, disjoint classes of possible predicaments. No one can correctly represent chip 17 both as Harry represents it and as Hermione represents it. But Harry's and Hermione's experiences are both veridical. Harry really does occupy a *six feet from chip 17* predicament, and Hermione really does occupy a *three feet from chip 17* predicament. So we've got both of their experiences coming out as veridical, despite their incompatibility. Since both of their experiences are veridical, we've obviously got no asymmetric error. Finally, we've got all of this because there's something in the neighborhood of projection going on – attribution of features like *being three feet away* to things, based not (just) on the objective features of the things perceived, but (also) on the peculiarities of

the perceiver and their particular perceptual situation. So in this sort of case, it looks as if we've satisfied all of INCOMPATIBILITY, NO ASYMMETRIC ERROR, CORRECTNESS, and PROJECTION.

Just as self-locating belief isn't restricted to geographical self-location, centering features aren't restricted to relative-position features like *being nearby* or *being three feet away*. Most importantly for our purposes, they include features that are tied up with the effects that things have on our sensory apparatus. This is particularly relevant because these sorts of subjective features look much more like the sorts of features that would help us cash out projectivism in the sorts of cases where it looks most appealing, and they point us toward my proposed projectivist account of the unique hues.

Let B be the phenomenal property that your experiences have when something looks unique blue to you. Here is a self-locating correctness condition: the one that's satisfied by all and only the predicaments the subjects of which are such that chip 17 is disposed to cause B sensations in them. (Making some not-too-contentious assumptions, this is just the set of predicaments whose subjects are disposed to have B sensations in response to chip 17.) That's the correctness condition that's determined by applying *unique blue* to chip 17.

Unique blue is a centering feature. When Hermione's visual experience attributes it to chip 17, it determines a self-locating correctness condition: the set S including all and only those predicaments whose inhabitants are disposed to get B sensations from chip 17 in normal viewing conditions. The experience is veridical iff its subject is in a predicament that's a member of S.

Equivalently, in centered-worlds terms: it determines the class of <w,t,i> triples such that at t, in w, i is disposed to get B sensations from chip 17 in normal viewing conditions. The experience is veridical at all and only those centered worlds. Equivalently, in property-self-

attribution terms: it determines a property that those who take such experiences at face value will self-attribute, and such that the experience will be veridical whenever it's had by someone with the property: *being disposed to get B sensations from chip 17 in normal viewing conditions*. The experience is veridical only if its subject has that property.[23]

Perceptual experiences that represent chip 17 as unique blue are veridical in all and only the *disposed to get B sensations from chip 17 in normal viewing conditions* predicaments. Everybody who thinks chip 17 is unique blue self-locates as someone who's disposed to get B sensations from chip 17 in normal viewing conditions.[24]

Again, this is not a case of *geographical* self-location in the way that attributions of *being nearby* or *being three feet away* are, but it is still, like those sorts of attributions, a case of a visual experience with self-locating correctness condition. That is, it's a case in which one's

[23] I'm simplifying a bit here. The reader is likely to have noticed that there are actually a number of different centering features available in this neighborhood, differing with respect to how one fixes the relevant circumstances. I chose the *normal circumstances* version for simplicity, but there's also, for any C that specifies some possible type of circumstance, the centering feature that, when attributed to an object x, determines the property *being disposed to get B sensations from x in C*. There's also, perhaps most interestingly, the centering feature that, when attributed to x, determines the property, *being disposed to get B sensations from x in present circumstances*. Attributing this centering feature to chip 17 determines (returning to centered worlds talk) the class of <w,t,i> triples such that i is disposed to get B sensations from chip 17 at t in w. I think that the *normal circumstances* versions are probably the best candidates to be the sorts of centering features that our ordinary perceptual attributions of unique hues are trading in, but I also think that this is a place where things are likely to get complicated. (See footnote 26 for one potential source of complication.)

[24] A couple words about how what I say here fits with things I've said elsewhere: It fits nicely with what I say in "Secondary Qualities and Self-Location" (Egan 2006a). The story here is, essentially, that we can cash out the distinction between the projected qualities and the rest in the same way that I want to cash out the distinction between the secondary qualities and the primary qualities in that paper. This shouldn't be terribly surprising – the things projectivists and secondary-quality theorists say, about how the qualities in question don't really reside in the things we attribute them to, aren't discovered parts of the objective structure of reality, etc., are pretty similar, and seem to call out for a similar sort of treatment. What I say here fits less well with what I say in "Appearance Properties?" (Egan 2006b). The story there is that, while the *colors* are full-fledged objective features of things in our environments, perceptual experience also trades in some other features that aren't like that. The story I tell here about unique hues is pretty much the story that I tell there about these additional, non-color qualities. So if what goes for unique hues goes for colors generally (which seems to me to be very plausible, but not quite obviously mandatory), what I say here and what I say there can't both be true. I do think that they're both plausible candidate views, and they're both views it's good to have on the table. Forced to choose one to endorse, I'm not certain which way I'd go.

visual experience has a correctness condition that specifies a set of possible predicaments, rather than a set of possible worlds. Here, the set of predicaments isn't unified by geographical or positional similarities, but by a certain response-dispositional similarity: they're not predicaments whose subjects are all in a similar sort of geographical location, but ones whose subjects all are disposed to respond in the same sort of way to the object in question.

Similarly, everybody who thinks chip 17 is greenish-blue self-locates as someone who's disposed to get a different, incompatible sort of sensation from chip 17 in normal viewing conditions (let's call them GB sensations). More generally, perceptual experiences that represent chip 17 as having some shade incompatible with unique blue will only be veridical in predicaments whose subjects are not disposed to get B sensations from chip 17 in normal viewing conditions. (Because they're disposed to get some other, incompatible sort of sensation – which one will depend on which incompatible shade is attributed.)

So Ron's and Hermione's visual experiences of chip 17 are incompatible – the correctness condition of Ron's experience is, in virtue of its attribution of *greenish-blue* to chip 17, disjoint from the correctness condition of Hermione's experience, in which she attributes *unique blue* to chip 17. Since there are no predicaments whose subjects are disposed to get both B and GB sensations from chip 17 in normal circumstances, the set of predicaments determined by attributions of *unique blue* to an object and the set determined by an attribution of *greenish-blue* to that object are going to be disjoint. And so it won't be possible for anybody to have a veridical experience that attributes both to the same object. So we've got INCOMPATIBILITY.

Both Ron's and Hermione's experiences are veridical. Ron really is in a *disposed to get GB sensations from chip 17 in normal viewing conditions* predicament. Hermione really is in a

disposed to get B sensations from chip 17 in normal viewing conditions predicament. So we've got CORRECTNESS. And since we've got CORRECTNESS, we've got NO ASYMMETRIC ERROR, because we've got no error at all. (That is, we've got no error at all in this case. Of course it's still possible for people, and their experiences, to be make erroneous attributions of these sorts of qualities – for example, when one is viewing the object in other-than-normal circumstances.)

This sort of account is able to deliver both CORRECTNESS and INCOMPATIBILITY because the contents in question are self-locating. The attributions of *unique blue* and *greenish-blue* to chip 17 are *incompatible* because no single individual can (simultaneously) correctly attribute both qualities to the same thing. The correctness conditions of the two attributions aren't jointly satisfiable (by a single predicament), because the classes of predicaments picked out are disjoint. The attributions are both *correct* because the people making the different attributions occupy relevantly different predicaments: Ron really is disposed to get GB sensations (and not B sensations) from chip 17, and Hermione really is disposed to get B sensations (and not GB sensations) from chip 17.

Finally, we've got PROJECTION: Part of how chip 17 is being represented in Ron's and Hermione's experiences is an artifact of contingent stuff about the perceivers' own particular perceptual endowments.

Other bits of how chip 17 is represented in Ron's and Hermione's experiences aren't like that – chip 17 is, e.g., represented as being square. That's on account of chip 17's squareness, and such experiences are veridical in only the *chip 17 is square* predicaments. (Equivalently: the predicaments located in *chip 17 is square* worlds.) If Ron and Hermione are worldmates, and one of them has an experience that represents chip 17 as square, while the other has an

experience that represents it as, say, circular, one of them has got to be wrong. So there's an interesting distinction between the projected bits of perceptual content (such as attributions of *unique blue*) and the non-projected bits of perceptual content (such as attributions of *squareness*).

This is an account on which there are, as the projectivist thought would have it, some bits of the contents of our perceptual experiences that aren't really tracking the objective features of objects in our environment. But this isn't because they aren't tracking anything at all. What the projected bits of perceptual content are tracking – by and large correctly – is aspects of the perceiver's particular predicament. More specifically, they're tracking something about her situation relative to the object that the projected features are attributed to. This is what allows for *projectivism* – some of the features attributed to objects in perception are attributed as a result not (just) of openness to what the objects in one's environment are like, but (also) on account of the peculiarities of one's own particular perceptual apparatus or situation – without error.

5. Objections and Responses

I'll close by briefly addressing two concerns, the responses to which will, I hope, help to clarify the proposal.

First concern: Maybe we're failing to respect some important distinctions here, by trying to give a one-size-fits-all account of a very diverse range of phenomena. This account lumps the traditional, paradigmatic candidates for projectivist treatment – colors and the like – into the same category with features like egocentric *nearness*. But these don't, on the face of it, belong in the same category. It seems as if there should be something distinctive about the colors,

which marks them off as different from these sorts of transparently self-locating egocentric-location features. A properly constructed projectivism should allow us to be projectivists about the colors while still respecting that difference.

More generally, we might be concerned that this sort of account will lump too wide a range of phenomena into a common category. Recall the diversity of the list of motivating phenomena we encountered earlier. We probably don't want a uniform account of all of those, which lumps perceptual affordances, unique hues, visual-field shapes, and so on all in one category. There are important differences between these different kinds of features, and we'll want our theory (or theories) of them to respect those differences.

But the sort of projectivist view I'm advocating isn't committed to the absence of any interesting distinctions between these sorts of perceptually attributed features. It's just committed to their having something important in common. This still allows for a lot of important diversity within the category of projected features.

For example: One important difference between location-in-egocentric-space features and the colors, on this account, will be that the first wear their involvement in self-locating representation on their sleeves in a way that the second don't.[25] Unlike the case of the colors, there's no temptation to treat egocentric-location features like *being nearby* and *being on the left* as fully objective, out-there-to-be-discovered properties of the objects in our environments. The self-locatingness of color attributions is an interesting theoretical discovery in a way that the self-locatingness of egocentric-location properties isn't. (Features like egocentric *nearness*, for

[25] At least, a difference between egocentric-location features and the unique hues or maximally specific shades. There's a bit of a step between projectivism for the unique hues and projectivism for colors in general, where one could potentially get off the bus.

example, don't even satisfy the requirements for what Joyce (forthcoming) calls "minimal projectivism", since we don't experience egocentric *nearness* as an objective feature of the world.)

Elsewhere in this volume, Jonathan Cohen discusses some reasons why it might not be immediately phenomenologically apparent that some perceptually attributed quality is *relational* rather than monadic. A very similar story can be told, I think, about why it wouldn't necessarily be obvious that a given feature attributed in perception is a centering feature rather than an objective property.

It could be, for example, that we're tempted to think of color properties as objective rather than self-locating because the telltale interpersonal variations in perception aren't immediately obvious, or because the relevant sorts of *intra*personal variation over time aren't as prevelant, or as salient, as the ones that we encounter in the case of, for example, locations in egocentric space.

It could also be (and here I am departing from Cohen, though I think that he could endorse a very similar thought) that part of the objective (and non-relational) appearance of some centering features is due to their lack of *explicitly* relational counterpart concepts. (Or, perhaps better, lacking readily available, explicitly relational counterpart concepts that are widely acknowledged, in common circulation, etc.) *Being three feet away* has readily available, commonly thought and talked about, explicitly relational counterparts which are themselves also represented in perception: *being three feet away from x* for various x. *Unique blue*, on the other hand, doesn't seem to have any such readily available, commonly talked and thought about, explicitly relational counterparts. We don't standardly talk or think about what's *unique blue for*

x for various x, nor does visual perception attribute these sorts of transparently relational properties to the objects in our environment. On the face of it, anyway, visual perception just takes a stand on what's *unique blue*, simpliciter. These sorts of cases – in which there isn't a readily available, explicitly relational counterpart – will plausibly be among the cases where it's most likely to be an interesting discovery that the content in question is self-locating, since they may not wear their relationality on their sleeves.

So one important internal division to draw within the category of centering features is, roughly, the line between the ones that wear their self-locatingness on their sleeves, and those whose implication in the determination of self-locating content is an interesting theoretical discovery. Let me also, by way of example, highlight a couple other potentially important internal divisions within the category.

One such division is the one between those centering features that have to do with relatively stable facts about the observer's perceptual apparatus, and those that have to do with relatively transient facts about her particular perceptual situation. This division will help us to draw some of the intuitive lines between the different motivating examples from early in the paper: visual-field shapes, for example, will fall on one side of the line, while the colors (on the sort of view outlined above) will fall on the other.[26] (To some extent, this distinction is likely to

[26] I think that, in fact, this sort of distinction also turns up between different sorts of color-features. I'm very sympathetic to Cohen's (this volume) distinction between the *steady* and *unsteady* colors. That's easy to capture on the sort of picture I'm offering here, by allowing for a distinction between two kinds of centering features we can attribute to objects in our environment: one kind that has to do with the responses that the objects are disposed to cause in us in *normal* circumstances, and another kind that has to do with the responses that the objects are disposed to cause in us in *present* circumstances. Attributions of the first sort of quality take a stand on what sorts of perceptual apparatus we have (and on the long-term, global features of our environment that determine what sorts of circumstances count as *normal*). Attributions of the second sort of quality also take a stand on the details of the particular, local, potentially fleeting details of our perceptual circumstances. (I'm also attracted to an explanation of color-constancy phenomena that exploits this distinction, where the differently illuminated bits of the wall are represented as having the same *steady* color, but different *unsteady* colors.)

follow the one just discussed – centering features having to do with the transient details of one's perceptual situation are, by and large, more likely to wear their self-locatingness on their sleeves than those having to do with the construction of one's perceptual apparatus. But the two won't track each other perfectly.)

Another important difference between the various centering features will be in the aspects of one's predicament that their attribution serves to characterize. Attributions of colors serve to characterize what we might, for lack of a better term, call one's *sensory circumstances* – how one's sensory apparatus is disposed to respond to particular objects in one's environment. Attributions of affordances, on the other hand – as when the hammer is represented as *to be used for hammering*, the nail as *to be hammered*, the cake as *to be eaten*, and the bubble wrap as *to be popped* – would serve, on this kind of view, to characterize one's *practical circumstances*, by representing the practical aspects of one's predicament.

Second concern: Does attributing self-locating contents to perceptual experience require overly fancy explicit self-representation? In particular, if attributing colors to things required explicit self-representation, that seems like it would be a problem. It sure seems as if creatures without the capacity for explicit self-representation could have color vision.

Response: No, it doesn't. It had better not – we need self-locating content to model cognition and perception in lots of creatures that lack the capacity for any kind of conceptually sophisticated explicit self-representation. The representational mental states of dogs and cats aren't any better modeled in terms of possible worlds than our own are. What's cognitively in common to Rex and Fido when they both want to go walkies, or when they both think that there's a squirrel nearby, is not going to be happily represented in possible-worlds terms. A

satisfactory theory of animal cognition, no less than human, requires the attribution of self-locating content.

What's required for a given mental or perceptual representation to have self-locating content is just for it to play the right sort of role in the host's cognitive and behavioral economy. Perceptual experiences have distinctive sorts of impacts on behavior, and tend to give rise to distinctive sorts of beliefs. The perceptual experiences with self-locating contents are the ones whose behavioral and cognitive effects are best modeled with self-locating propositions rather than possible-worlds propositions. For example, a visual experience that disposes the creature that has it to behave in ways that would be appropriate/adaptive/whatever in a certain distinctive sort of predicament, rather than a certain distinctive sort of *world*. (Ones whose content is better modeled as, for example, *Tiger nearby!* than *Tiger near Fred!*. The first sort of content will be a good candidate to be what's in common to every prey animal's tiger-in-immediate-foreground sorts of visual experiences, well suited to explain the common behavioral upshots of such experiences in the various animals that have it. The second not so much.)

Conclusion

I think that there's a lot to be said for an account according to which perception trades both in centering features and objective properties, both as a way of cashing out the projectivist thought and as part of the correct theory of perception. I think it's an interesting question just which perceptually attributed features should get which sort of treatment, though I think there's a lot to be said for the view that the unique hues, and colors generally, fall on the centering features side of the divide. Making this distinction between the kinds of features attributed in perception is, I think, the most attractive way being a projectivist without being an error theorist.

Allen, Keith. 2007. The Mind Independence of Colour. *European Journal of Philosophy* 15: 137-158.
Blackburn, Simon. 1984. *Spreading the Word: Groundings in the Philosophy of Language*, Oxford: Clarendon Press.
Block, Ned. 1999. Sexism, Racism, Ageism and the Nature of Consciousness. *Philosophical Topics* 26: 71-88.
Bradley, Peter and Tye, Michael. 2001. Of Colors, Kestrels, Caterpillars, and Leaves. *Journal of Philosophy* 98: 469-487.
Byrne, Alex and Hilbert, David. 2007. Truest Blue. *Analysis* 67: 87-92.
Castañeda, Hector-Neri. 1966. He: A Study in the Logic of Self-Consciousness. *Ratio* 7: 130-57.
Castañeda, Hector-Neri. 1967. Indicators and Quasi-Indicators. *American Philosophical Quarterly* 4: 85-100.
Castañeda, Hector-Neri. 1968. On the Logic of Attributions of Self-Knowledge to Others. *Journal of Philosophy* 65: 439-56.
Chisholm, Roderick. 1982. *The First Person: An Essay on Reference and Intentionality*. Minneapolis: University of Minnesota Press.
Cohen, Jonathan. 2004. Color Properties and Color Ascriptions: A Relationalist Manifesto. *The Philosophical Review* 113: 451-506.
Cohen, Jonathan. 2006. Color and Perceptual Variation Revisited. *Dialectica* 60: 307-319.
Cohen, Jonathan. 2007. A Relationalist's Guide to Error about Color Perception. *Nous* 41: 335-353.
Cohen, Jonathan. this volume. Do Colors Look Relational?
Cohen, Jonthan. Forthcoming. *The Red and the Real: An Essay on Color Ontology.* Oxford: Clarendon Press.
Cohen, Jonathan, Hardin, C.L., and McLaughlin, Brian P. 2006. True Colors. *Analysis* 66: 335-340.
Cohen, Jonathan, Hardin, C.L., and McLaughlin, Brian P. 2007. The Truth about 'The Truth about True Blue'. *Analysis* 67: 162-6.
Egan, Andy. 2006a. Secondary Qualities and Self-Location. *Philosophy and Phenomenological Research* 72: 97-119.
Egan, Andy. 2006b. Appearance Properties? *Noûs* 40: 495-521.
Egan, Andy. Forthcoming. Disputing about Taste. In *Disagreement*, ed. Richard Feldman and Ted Warfield. Oxford: Oxford University Press.
Egan, Andy. MS. Three Grades of Self-Involvement: Self-Locating Content in Thought and Language.
Gibson, James. 1977. The Theory of Affordances. In *Perceiving, Acting and Knowing*, ed. Robert Shaw and John Bransford. John Wiley & Sons.
Gibson, James. 1979. *The Ecological Approach to Visual Perception*. Boston: Houghton-Mifflin
Hardin, C.L. 1988. *Color for Philosophers: Unweaving the Rainbow*. Indianapolis: Hackett.
Heidegger, Martin. 1927/1962. *Being and Time*. New York: Harper & Row.
Hume, David. 1740/1978. *A Treatise of Human Nature*. Ed. L.A. Selby-Bigge. Oxford: Clarendon Press.

Hume, David. 1751/1983. *An Enquiry Concerning the Principles of Morals*. Cambridge, MA: Hackett.
Hume, David. 1757/1965. Of the Standard of Taste. In *Of the Standard of Taste and Other Essays*, ed. John W. Lenz. Indianapolis: Bobbs-Merrill.
Jackson, Frank and Pargetter, Robert. 1987. An Objectivist's Guide to Subjectivism about Colour. *Revue Internationale de Philosophie* 41:127-41.
James, William. 1907. *Pragmatism*. New York: Longmans, Green & Co.
Johnston, Mark. 1992. How to Speak of the Colors. *Philosophical Studies* 68: 221-263.
Joyce, Richard. Forthcoming. Is Moral Projectivism Empirically Tractable? Forthcoming in *Ethical Theory and Moral Practice*.
Kalderon, Mark. 2007. Color Pluralism. *Philosophical Review* 116: 563-601.
Kelly, Sean. this volume. The Normative Nature of Perceptual Experience
Lewis, David. 1979).Attitudes *De Dicto* and *De Se*. *Philosophical Review* 88: 513-43. Reprinted with postscripts in David Lewis, *Philosphical Papers v.1*. Oxford: Oxford University Press.
Locke, John. 1690/1959. *An Essay Concerning Human Understanding*. New York: Dover.
Matthen, Mohan. this volume. How Things Look (and What Things Look That Way)
McLaughlin, Brian. 2003. The Place of Color in Nature. In *Colour Perception: Mind and the Physical World*, ed. Rainer Mausfeld and Dieter Heyer. New York: Oxford University Press.
Millner, A. David and Goodale, Melvin A. 2008. Two Visual Systems Re-Viewed. *Neuropsycholgia* 46: 774-785.
Mizrahi, Vivian. 2006. Color Objectivism and Color Pluralism. *Dialectica* 60: 283-306.
Noë, Alva. 2004. *Action in Perception*. Cambridge, MA: MIT Press.
Peacocke, Christopher. 1992. *A Study of Concepts*. Cambridge: MIT Press.
Perry, John. 1979. The Problem of the Essential Indexical. *Nous* 13:3-21.
Quine, W.V. 1969. Propositional Objects. In *Ontological Relativity and Other Essays*. New York: Columbia University Press.
Rosen, Gideon. 1994. Objectivity and Modern Idealism: What is the Question?. In *Meaning in Mind*, ed. Michaelis Michael and John O'Leary-Hawthorne. Dordrecht: Kluwer.
Schellenberg, Susanna. 2007. Action and Self-Location in Perception. *Mind* 116: 603-632.
Schellenberg, Susanna. Forthcoming. The Situation-Dependency of Perception. Forthcoming in *Journal of Philosophy*.
Shoemaker, Sydney. 1990. Qualities and Qualia: What's in the Mind? *Philosophy and Phenomenological Research* 50 (Supp):109-131.
Shoemaker, Sydney. 2006. On the Ways Things Appear. In *Perceptual Experience*, ed. Tamar Gendler Szabo and John Hawthorne.
Siegel, Susanna. (this volume) Do Visual Experiences Have Contents?
Sosa, Ernest. 1983. Propositions and Indexical Attitudes. In *On Believing: Epistemological and Semiotic Approaches*, ed. Herman Parret. Berlin: Walter de Gruyter & Co.
Stalnaker, Robert. 1981. Indexical Belief. *Synthese* 49:129-151.
Thompson, Brad. 2007. Shoemaker on Phenomenal Content, *Philosophical Studies* 135:307-334.
Thompson, Brad. Forthcoming. Sense for Senses, *Australasian Journal of Philosophy*.
Tye, Michael. 2006a. The Puzzle of True Blue. *Analysis* 66: 173-178.

Tye, Michael. 2006b. The Truth about True Blue. *Analysis* 66 (292): 340-344.
Tye, Michael. 2007. True Blue Redux. *Analysis* 67 (293):92-93.