

Appearance Properties?¹

Introduction

Intentionalism—the view that the phenomenal character of an experience is wholly determined by its representational content—is very attractive. Unfortunately, it’s in conflict with some quite robust intuitions about the possibility of phenomenal spectrum inversion without misrepresentation. Faced with such a problem, there are the usual three options—reject intentionalism, discount the intuitions and deny that spectrum inversion without misrepresentation is possible, or find a way to reconcile the two by dissolving the apparent conflict.

Sydney Shoemaker’s (1994) introduction of *appearance properties* is a particularly ingenious way of pursuing the third strategy, by maintaining that there is a representational difference between the phenomenally spectrum-inverted subjects.² In introducing appearance properties, Shoemaker does two things: he identifies a theoretical role for some family of properties to play, and he suggests a family of properties as candidates to play that role. I’ll argue that his proposed candidates do not play the role as well as we would like, suggest some new candidates, and argue that they do a better job

¹ Thanks to Alex Byrne, Tyler Doggett, Stephen Yablo, Ned Hall, Robert Stalnaker, James John, Uriah Kriegel, David Chalmers, John O’Dea, Daniel Stoljar, Philippe Chuard, and the MATTI group for extremely helpful discussions, comments and suggestions.

² Shoemaker (1994, 2000, 2001). In these papers he calls them ‘phenomenal properties’. He’s recently changed his terminology (in Shoemaker (2003, forthcoming)), and I follow him in the new terminology here.

of playing the role. The reason for the question mark in the title is that, if I'm right, it turns out that the best candidates to play the appearance property role aren't properties.

1. Intentionalism and Spectrum Inversion

Intentionalists hold that the phenomenal character of an intentional mental state supervenes on its representational content. Visual experiences are intentional mental states.³ So, if intentionalism is true, there can't be two visual experiences that are alike in their representational content, but differ in their phenomenal character.

Suppose Ernie and Vert are phenomenally spectrum inverted with respect to each other. Ernie's visual experiences when he looks at ripe tomatoes, fire engines, and cooked lobsters are phenomenologically just like ours.⁴ Vert's visual experiences are phenomenally inverted. His experiences when he looks at ripe tomatoes, fire engines, and cooked lobsters are phenomenally like our (and Ernie's) visual experiences when we look at unripe tomatoes, limes, and uncooked lobsters. We can put this in terms of the qualitative features of Ernie's and Vert's experiences by saying that Ernie's visual experience of Kermit is G, while Vert's is R, where a G experience is one with the phenomenal character of our experiences of unripe tomatoes, etc., and an R experience is one with the phenomenal character of our experiences of ripe tomatoes, etc.

³ Or at least, being the subject of a particular visual experience is an intentional mental state. If there is an important state/experience distinction, nothing bad will come of my playing a bit fast and loose with it here.

⁴ Let's assume for simplicity that there isn't any *actual* spectrum inversion, so that our visual experiences really are the same.

Suppose also that Ernie and Vert are (in relevant respects) historically and behaviorally indistinguishable. They've both had relevantly similar sorts of causal intercourse with limes and lobsters, they both use color terms in the same way, they both put things into the same piles when asked to sort them by color, and so on. For example, when Ernie and Vert look at Kermit, they both say "he's green", and they both stack Kermit in the pile with the unripe tomatoes and uncooked lobsters when asked to sort things by color.

A natural interpretation of Ernie and Vert's situation is that it's one in which Ernie and Vert are phenomenally spectrum inverted, but their visual experiences agree on the colors of things. If their experiences *didn't* agree on the colors of things, one or the other of them would have to be getting the colors of things *wrong*. And this seems implausible. Given the similarity of their histories, discriminatory abilities, stacking behavior, etc., we don't seem to have any reasonable grounds for attributing the error to one of them rather than the other. (This is where the 'without misrepresentation' part of 'spectrum inversion without misrepresentation' comes in.) So we ought to say that their visual experiences both represent Kermit as being green, even though their phenomenal character is different.

Let's suppose that this kind of case—in which two observers are phenomenally inverted, though their visual experiences have the same color content—is possible.⁵ Then

⁵ For some arguments that it's not, see e.g. Byrne and Hilbert (1997), Harman (1990), Hilbert and Kalderon (2000), and Stalnaker (2000). Note that if Stalnaker is right, then it's not just that cases like this aren't possible—it's that it doesn't even make sense to *talk* about cases like this.

we seem to have a counterexample to intentionalism. That is, we seem to have a difference in phenomenal character without a difference in representational content.

However, an important part of the argument was suppressed in the previous paragraphs. Why should we agree that Ernie's and Vert's visual experiences have the same (overall) representational content? Here is a plausible argument: if Ernie's and Vert's experiences don't differ with respect to *color* content, they don't differ with respect to representational content *at all* (or at least, not in any way that's potentially relevant). The only available place to locate a representational difference between Ernie's and Vert's experiences is in which colors their visual experiences represent Kermit as having. So since Ernie's and Vert's experiences have the same color content (both represent Kermit as being green), they have the same representational content *simpliciter*.

Shoemaker's strategy is to deny the claim that Ernie's and Vert's experiences differ with respect to representational content *only* if they differ with respect to color content. Intentionalism is to be saved by supposing that, while Ernie's and Vert's experiences both represent Kermit as being *green*, there is a representational difference with respect to the *other* properties that they attribute to Kermit.

2. *A Bunch of Distinctions, a Problem, and Another Distinction*

Before moving on to Shoemaker's proposal, we should pause to head off an alternative line of response to the conflict between intentionalism and the possibility of spectrum inversion without error.

Contentful experiences have *representational properties*. One kind of representational property is the property of having a certain content. There are a lot of competing ideas about what contents are, but all parties to the debate must concede that representations have *at least* the following sort of content: they make a distinction between the possibilities in which things are as they're represented to be, and the possibilities in which things are otherwise. Maybe there are other, finer-grained kinds of contents, too, but if there are, those kinds of finer-grained contents will each at least *determine* a coarse-grained possibility-carving content of the sort described above.

So one (fairly coarse-grained) class of representational property is the class of possibility-carving properties. Call these possibility-carving properties the *pure* representational properties.⁶ (I will, from now on, stipulatively use 'content' to mean 'possibility-carving content'. I don't mean to commit myself to any view about which notions of content are viable, theoretically useful, etc. by making this stipulation—I just need a less cumbersome expression than 'possibility-carving content', and I won't be employing any of the competing notions in what follows.)

When Grover believes, Oscar asserts, Miss Piggy fears, and Ernie and Vert see that Kermit is a frog, Grover's belief, Oscar's assertion, Miss Piggy's fear, and Ernie's

⁶ For now, it does no harm to read 'possibility-carving content' as 'possible-worlds content'. I'm hedging now, though, because the difference will matter later on.

and Vert's visual experiences all share a pure representational property. They all have contents that separate the worlds in which Kermit is a frog from the rest.

There are other kinds of representational properties. For example, *visually representing that Kermit is a frog*, or *representing Kermit, under mode of presentation K, as a frog*. Call these and their ilk *impure* representational properties. We can think of the pure representational properties as very general properties of *representing content C somehow-or-other*, and the impure representational properties as properties of *representing content C this way*.⁷

The intentionalist slogan is, "same representational properties, same phenomenal character". The distinction between pure and impure representational properties lets us draw a distinction between two sorts of intentionalism. *Pure intentionalism* says that the phenomenal character of experience supervenes on the pure representational properties of experience. *Impure intentionalism* says only that phenomenal character supervenes on impure representational properties.

Let's also distinguish between *maximal* and *submaximal* representational properties of both kinds. Maximal representational properties specify the content (or the content plus the manner of representation) of an experience exhaustively, while submaximal representational properties specify only part of the content of the

⁷ This more or less follows Chalmers' distinction between pure and impure representational properties in Chalmers (2004). The notion of content I've employed is more restrictive than his, though. It's actually a bit more restrictive than I need for the points I'm about to make, but it's difficult to state the weaker notion in a happy way.

experience.⁸ So pure intentionalism says that maximal pure representational properties fix phenomenal character, while impure intentionalism makes the weaker claim that maximal *impure* representational properties fix phenomenal character.

If it's possible to have both an occurrent thought (or some other sort of cognitive experience) and a visual experience with the same possibility carving content—that is, the same maximal pure representational property—then pure intentionalism is false. (Since while the visual experience will have some distinctive phenomenal character, having the thought either won't have any phenomenal character at all, or will have a very different one than that of the visual experience with the same content). Since this is almost certainly possible, pure intentionalism is almost certainly false. Call this the *problem of common content*. So it looks as if the intentionalist will have to appeal to impure representational properties in order to make her thesis plausible, given that there are some maximal pure representational properties that can be shared by experiences in different modalities (visual, cognitive, auditory, etc.).⁹

So why be so worried about the sort of spectrum inversion case described above? What the case of Ernie and Vert is a clear counterexample to is *pure* intentionalism, and

⁸ More carefully, and more generally: for any kind of property K, a property F is a maximal K property iff, for any K property G, something's having F either entails that it has G or entails that it lacks G. That is, F is a maximal K property iff something's having F fixes *all* of its K properties.

⁹ Alex Byrne (2001) discusses this sort of problem, and makes what I take to be the same distinction between varieties of intentionalism (though under a slightly different mode of presentation). He also offers some reasons for resisting this argument against pure intentionalism, so things are probably not so clear as I've suggested. (For one thing, while it's clear that experiences in different modalities can share very many *submaximal* representational properties, it's actually much less clear that they can share *maximal* representational properties). Thanks to James John and Alex Byrne for discussion. See John (2004) and Byrne (2001) for more details.

we already knew that that was false, because of the problem of common content. For all that's been said so far, it could be that Ernie's and Vert's experiences have the same possibility-carving content, but they each represent it in different ways. They might, for example, each represent the property *being green* under a different mode of presentation.¹⁰ So Ernie's and Vert's experiences would share all of their pure representational properties, while differing in their impure representational properties. Since the intentionalist already needs to appeal to impure representational properties to handle the problem of common content, where's the harm in appealing to them here as well?

Notice a distinction between two kinds of impure representational properties. First, there are impure representational properties like *visually representing that Fozzie is on stage*, *auditorily representing that there's a rooster nearby*, and *representing in belief that no armadillo weighs more than any aircraft carrier*. These are impure because they specify the sensory or cognitive modality in which the content is being represented. Call these representational properties *modality-impure*. Everybody has to concede that our experiences have modality-impure representational properties—it's just a plain fact that there are a number of different representational modalities. (Not everybody has to concede that these properties are theoretically interesting, but everybody *does* have to concede that our experiences have them.)

¹⁰ See Chalmers (2004) for such an account.

Compare these to impure representational properties like, *represents that Fozzie is on stage (where Fozzie is presented under mode C)*, *represents that Fozzie is on stage (where Fozzie is presented under mode S)*, and *represents that Kermit is green (where being green is presented under mode G)*. Call these representational properties, which are impure because they specify modes of presentation for some of the things represented, *mode-of-presentation impure*, or *MOP-impure*. Not everyone will want to concede that our experiences have MOP-impure representational properties. It's a controversial bit of theory that there are different modes of presentation under which we represent things (and properties). So while appealing to modality-impure representational properties to rescue intentionalism from the problem of common content is relatively theoretically innocent, appealing to MOP-impure representational properties requires some substantial theoretical commitments.

An intentionalism that allows for differences in phenomenal character due to differences in *modality-impure* representational properties, but not due to *MOP-impure* representational properties, is a close relative of pure intentionalism. It's a theory on which pure intentionalism is true *within each modality*, though not across modalities. Call such a theory a *modality-impure intentionalist* theory. Within a given representational modality, phenomenal character supervenes on *pure* representational properties—no two experiences in the same modality can differ in their phenomenal

character without differing in their pure representational properties.¹¹ This is the sort of intentionalist theory that I'll be concerned to square with the possibility of spectrum inversion without error.

Modality-impure intentionalism is an attractive view because it seems to be the version of intentionalism with the fewest possible theoretical commitments. This sort of intentionalist theory will put fewer constraints on what we're allowed to say in the rest of our theorizing than a version that, for example, appeals to MOP-impure representational properties as well.¹² Cases of spectrum inversion without misrepresentation are *prima facie* counterexamples to such a theory, since Ernie and Vert are both having *visual* experiences with the same (possibility-carving) content. Shoemaker's introduction of appearance properties is a promising line of response to these apparent counterexamples.

3. Appearance Properties

The theoretical motivation for Shoemaker's move is pretty clear: adding an extra layer of content to visual experience allows us to agree that Ernie's and Vert's experiences are the same with respect to color content, while denying that they're the same with respect to representational content *simpliciter*. So we can be intentionalists

¹¹ This distinction is the same as Byrne's (2001) distinction between *intermodal* and *intramodal* intentionalism.

¹² Though one thing that might happen is that (a) it turns out that an intentionalism that appeals to MOP-impure representational properties is much less messy than one that doesn't, and (b) we're unavoidably committed to modes of presentation anyway, for independent reasons. If that's how things shake out, then we'll probably (assuming that we want to be intentionalists of some stripe) want to go with the MOP-impure version of intentionalism. But it's not at all clear that that *is* how things will shake out, and anyway, there's an interesting chunk of theoretical space to get mapped out here—what's the best story for a modality-impure intentionalist to tell about spectrum inversion?

(modality-impure intentionalists, that is—this qualification is hereafter omitted) while allowing for the possibility of spectrum inversion without misrepresentation.

The addition of an extra layer of content also has some appeal at a less theoretical level. The intuitive idea behind the proposal is that, while Ernie's and Vert's experiences agree about what *color* Kermit is, they still, in *some* sense or other, represent him as *looking different*. So there should be some range of properties that aren't colors, but still mark differences in how things look, with respect to which Ernie's visual experience is different from Vert's.¹³

This identifies a role that we'd like to have some properties to play—marking a difference in how things (representationally) look to Ernie and Vert that isn't a difference in which colors their visual experience attributes to things. Call whichever properties (if any) play this role the *appearance properties*, and call the one that Ernie represents Kermit as having 'PG' (for 'phenomenal green') and the one Vert represents him as having 'PR' (for 'phenomenal red').¹⁴

Now the proponent of appearance properties has an important question to answer—which properties are the appearance properties, exactly? For our purposes, we

¹³ This isn't, by itself, a sufficient motivation for positing some further sort of content with respect to which Ernie's and Vert's experiences differ. Given that there is (at least in philosophical parlance) a distinction between *intentional* and *phenomenal* senses of 'seems' (one that Shoemaker (1982) either helpfully points out or perniciously introduces, depending on your view of the distinction), we might well want to say that there's a similar distinction between intentional and phenomenal senses of 'looks', and that things look different to them only in the phenomenal sense. I think, though, that this sort of objection is largely headed off by Shoemaker's more recent (Shoemaker (2003, forthcoming) appeal to, for example, cases in which some object seems to have a uniform color, but parts of it look different because some are brightly lit and others are in shadow, as another source of motivation for the introduction of appearance properties.

¹⁴ These names are due to Shoemaker's old terminology, on which the appearance properties were called *phenomenal properties*.

can treat this as the question, *which properties are PG and PR?* We should be skeptical about the prospects of the appeal to appearance properties until we've been given some idea of what these properties *are*, exactly.¹⁵

Let's look at the role that the appearance properties are supposed to play in some more detail. There are some constraints on what the appearance properties can be like, if they're going to allow us to reconcile intentionalism with the possibility of spectrum inversion without misrepresentation. For ease of presentation, I'll state them as constraints on what our two particular appearance properties—PG and PR—can be like. The generalizations should be clear enough.

DIFFERENCE: Ernie's and Vert's experiences have to represent Kermit as having *different* appearance properties. So it has to turn out that $PG \neq PR$.

It's pretty clear why we need this. What we need in order to reconcile the phenomenal difference between Ernie's and Vert's experiences with the intentionalist's

¹⁵ Is it so clear that we need to be given an informative account of which properties are the appearance properties? Stephen Yablo (1995) has argued that sometimes we ought to resist these sorts of demands for alternative specifications of which property F is—sometimes the specification that the objector wants an alternative to is the best one that's available. This seems pretty plausible for already well-entrenched properties, but not so good for newly introduced, theoretically motivated properties. Plausibly, we're not obliged to say which property *being red* is in any more informative terms. If, on the other hand, I introduce some property as *whichever property plays fancy theoretical role R*, then I think that I *am* under some obligation to show that there really is some antecedently respectable property that can play role R. In particular, in the case where I'm introducing a property as the one represented in a certain kind of experience, I'll need to say something at least a little bit informative about how things are represented as being in those experiences, and this will require saying something about the nature of the represented property.

supervenience claim is an accompanying representational difference. Adding more content to their visual experiences won't help unless the content we add to Ernie's experience is *different* from the content we add to Vert's.

CORRECTNESS: Ernie's and Vert's visual experiences should both be representing Kermit *correctly* when they represent him as being, respectively, PG and PR.

There are two motivations for this constraint. First, we're trying to find a way to reconcile intentionalism with the possibility of spectrum inversion *without misrepresentation*, and it would be strange to do this by simply relocating the misrepresentation—it would be strange to avoid having to attribute misrepresentation of *color* properties by replacing it with misrepresentation of *appearance* properties. (Though maybe this is not *as* bad—maybe it's better to have systematic error about weird, esoteric properties than to have systematic error about ordinary, everyday properties like colors.)

Second, and more persuasively, denying CORRECTNESS will force us to say that either Ernie's or Vert's visual experience, or both, systematically misrepresents the world (since one or both will represent things as having appearance properties that they don't have). If we say this, our choice of which one to attribute the error to will be objectionably arbitrary—there's no relevant difference between them that could ground a principled choice of one rather than the other as the one that's getting things right.

Saying that they're *both* getting things wrong requires us to attribute a lot of systematic misrepresentation (not just to Ernie and Vert, but in our own case, as well) and this ought to be avoided if possible.¹⁶

It's worth noticing a requirement that follows from CORRECTNESS and the constraint that Ernie's and Vert's visual experiences agree on their color content:

NOVELTY: The appearance properties aren't the colors. In particular, PG isn't *being green* and PR isn't *being red*.

(This is worth noticing because some of Shoemaker's candidates wind up looking suspiciously color-like.)

The above requirements are non-negotiable. No properties that fail to satisfy them can be the appearance properties, because they can't play the right role in reconciling intentionalism with the possibility of spectrum inversion, which is the whole reason for wanting to have appearance properties around in the first place.

There are three more features that it would be nice to have in the appearance properties, because they're intuitively well motivated, but which aren't absolutely required in order to do the work appearance properties were introduced to do.

¹⁶ Though see Boghossian & Velleman (1989) for dissent. Thau (2001) leaves it open that visual experience *might* systematically misrepresent things as having a special range of properties that they don't actually have. See Shoemaker (forthcoming) for a related discussion of, and concerns about, this feature of Thau's view.

SAMENESS: The appearance property that Ernie's visual experience attributes to Kermit is the same as the appearance property that Vert's visual experience attributes to, for example, cooked lobsters and ripe tomatoes.¹⁷

The idea here is that phenomenal spectrum inversion should go hand in hand with a certain kind of *representational* inversion. The reason why Ernie's and Vert's visual experiences are *phenomenally* inverted is that they are *representationally* inverted with respect to certain appearance properties.¹⁸

So if Ernie's visual experience represents paradigm green things as being PG and Vert's represents them as PR, Ernie's visual experience should represent paradigm red things as PR and Vert's should represent them as PG. The appearance property whose representation underlies the phenomenal character of Ernie's experience of green things (like Kermit) should be the same as the appearance property whose representation underlies the phenomenal character of Vert's experience of red things (like ripe tomatoes).

¹⁷ Don't be thrown off by the presence of both a DIFFERENCE and a SAMENESS requirement—I'm not imposing contradictory demands on the appearance properties. (At least, not yet.) DIFFERENCE requires that the appearance property that Ernie's visual experience attributes to Kermit be different from the one Vert's experience attributes to Kermit. SAMENESS requires that the appearance property that Ernie's visual experience attributes to Kermit be the same as the one that Vert's experience attributes to ripe tomatoes.

¹⁸ Another way to put the requirement: it's not just that having a certain maximal representational property necessitates that an experience has a certain maximal qualitative property, but that some submaximal representational properties (like *representing something as PG*), are such that any experience that has them also has some (submaximal) qualitative property.

The intuitive idea behind appearance properties is that they're the properties, representation of which underlies, explains, or maybe just *is*, the phenomenal character of our experiences. If we want to hang on to *that* claim, then we'll want to have SAMENESS. Then we can say that the reason why Ernie's and Vert's *phenomenology* is inverted is because the appearance property content of their visual experiences is inverted. The reason why Ernie's experiences of ripe tomatoes are like Vert's experiences of unripe ones is because the (salient, phenomenal-character-determining) property that Ernie's visual experience attributes to ripe tomatoes is the same one that Vert's experience attributes to unripe ones. Without SAMENESS, we don't have this explanation available.

CONTRARINESS: PG and PR should be *contraries*—correctly representing something as PG should be incompatible with correctly representing it as PR.

When Ernie learns that Kermit looks *that* way, he should learn that Kermit *doesn't* look (in the relevant respect) the way a ripe tomato looks, or the way Kermit looks to Vert. Mutatis mutandis for Vert.

There's an obvious problem with CONTRARINESS—it seems to be flatly inconsistent with CORRECTNESS. If Ernie's experience represents Kermit as PG, and Vert's represents him as PR, and they're both correct, then Kermit *is* both PG and PR, so PG and PR can't be incompatible in the way that CONTRARINESS seems to require. So

it seems a bit odd to look for a proposal that gives us CONTRARINESS, given that it's inconsistent with one of our non-negotiable desiderata.

Here are three things to say in response to this concern about whether or not we ought to be at all concerned about trying to satisfy CONTRARINESS: First, it turns out that the requirements of CONTRARINESS and CORRECTNESS *can* be reconciled, if we adopt the right view about what kinds of features are attributed to things in experience. (I'll explain how to effect the reconciliation in sections 5-6.) Second, Shoemaker seems to be right about the intuitive attractiveness of CONTRARINESS, and so a proposal that preserved it would be better than one that gave it up. Finally, given SAMENESS, there's a principled theoretical motivation for wanting CONTRARINESS.

There are some incompatibilities at the level of phenomenal character. The distinctive phenomenal feature of experience that goes along with representing something as PG is incompatible with the distinctive phenomenal feature of experience that goes along with representing that thing as PR.¹⁹ So given SAMENESS, it looks like it's impossible to have a visual experience that, for example, represents Kermit as both PG and PR—there's a certain sort of representational content that's not available to us (or at least, not available in visual experience). If PG and PR are contraries—if nothing can *be* both PG and PR—then it's not very surprising that we can't *represent* anything as both PG and PR. If PG and PR aren't contraries—if things can be both PG and PR, but we

¹⁹ The incompatibility isn't between G and R, but between some very much more specific, extremely difficult to describe phenomenal properties. There may be a horrible problem lurking in this neighborhood, but if there is, it's one that I won't be able to discuss here. So I pass over it in near-silence, hoping that no one will make a fuss.

can't represent them that way—then we've got a brute restriction on which of the perfectly possible states of affairs (of a kind that visual experience is in the business of representing) we're capable of representing in visual experience. This seems to be exactly the sort of consideration that makes us so sure that *being red* and *being green* are contraries. It would be nice to have CONTRARINESS in order to avoid this sort of brute restriction on our representational capacities.^{20,21}

Another desideratum that Shoemaker (1994) proposes is:

CONSTANCY: The appearance properties should be features that are had by things even when unobserved.

This is certainly the least well motivated of the desiderata. For one thing, it's not so clear how much of a pretheoretical, intuitive picture of the nature of appearance properties we really have, and we need a pretty detailed one to motivate this desideratum. For another, it's not clear how much any picture we *do* have ought to constrain what we say about properties that are so theoretically motivated and distant from ordinary talk.

²⁰ It's reasonable to worry about how much we ought to trust the inference from our inability to represent a situation in which P in visual experience to the impossibility of P. It's pretty clearly *not* a reliable inference across the board. We can avoid a lot of the apparent counterexamples by restricting our attention to the sorts of situations that visual experience is *in the business* of representing (the colors and shapes of medium-sized physical objects, for example, but not much of anything about the very small or the abstract, and not the densities, origins, or futures even of the medium-sizers). We also ought to make sure that we don't claim (even in the good cases) an *entailment* between inability to represent and impossibility.

²¹ Thau (2001) makes a similar appeal to CONTRARINESS, or a principle very like it, in criticizing Shoemaker's view.

In fact, in more recent work, Shoemaker appeals to two kinds of appearance properties, only one of which satisfies CONSTANCY. But still, CONSTANCY does sound pretty plausible on its face—I think that the intuition Shoemaker appeals to is there, even if tentative and of uncertain authority.

4. Finding the Appearance Properties

To make the sort of move Shoemaker is proposing work, we need to find some candidates to be the appearance properties. The successful candidate will be the one (if any) that does the best job of satisfying all of the desiderata, out of the candidates that satisfy them at least minimally well. It will have to satisfy all of the non-negotiable desiderata—DIFFERENCE (PG and PR have to be distinct), NOVELTY (the appearance properties can't be the colors), and CORRECTNESS (neither Ernie nor Vert is misrepresenting Kermit). It will also have to do a better job than the other candidates (that satisfy the non-negotiable desiderata) of satisfying the 'bonus' desiderata: SAMENESS (the property that Ernie attributes to Kermit is the one that Vert attributes to cooked lobsters), CONSTANCY (things can have appearance properties while unobserved) and CONTRARINESS (when Ernie learns that Kermit is PG, he learns that Kermit is not PR).

Shoemaker considers two candidates to be the appearance properties. One of them satisfies the minimal desiderata and SAMENESS, but satisfies neither CONSTANCY nor CONTRARINESS. The other appears, at a first pass, to satisfy the

minimal desiderata and SAMENESS, plus CONSTANCY as well. This seemingly better candidate, however, is actually ambiguous between three different proposals. On closer examination, two of these fail to meet the minimal desiderata, and the third delivers an implausible result about the appearance properties of non-actual things (and also fails to respect the intuition that motivates CONSTANCY). None satisfy CONTRARINESS, and all have some difficulty with NOVELTY.

The first of Shoemaker's candidates is the family of properties such as *currently producing a G experience in some observer*. (Recall that a G experience is one with the phenomenal character of our, and Ernie's, experiences of unripe tomatoes, limes, and uncooked lobsters, and an R experience is one with the phenomenal character of our experiences of ripe tomatoes, fire trucks, and cooked lobsters.)²² Call these, following Shoemaker, the *occurrent* appearance properties.

These properties satisfy DIFFERENCE, CORRECTNESS, NOVELTY, and SAMENESS, but not CONSTANCY or CONTRARINESS. They satisfy DIFFERENCE because something can be causing a G experience but not an R experience, and vice versa, so something can be PG without being PR (and vice versa). They satisfy CORRECTNESS because Kermit is causing a G experience in Ernie and an R experience in Vert, and so he really is both PG and PR. NOVELTY is satisfied, since it's clear that

²² Note that this won't give us a *reductive* intentionalism, since the representational properties are characterized in terms of phenomenal properties of experience. A nonreductive intentionalism is still worth pursuing, though. The question of the relation between representational content and phenomenal character is interesting independent of its connections to the mind-body debate, and the view that the connections are very, very close is quite attractive, independent of its consequences elsewhere.

the colors aren't these kinds of properties—the colors of things don't depend on their being presently perceived. Finally, they satisfy SAMENESS because it's clear that the property Ernie's visual experience attributes to Kermit is the one that Vert's attributes to ripe tomatoes, and vice versa.

These occurrent appearance properties pretty clearly don't satisfy CONSTANCY because Kermit won't be producing any G or R experiences in observers when he's not being observed, and they don't satisfy CONTRARINESS because Kermit *is* both PG and PR, which shows that they're not contraries. Shoemaker (2000) dismisses these properties in favor of the next candidates because of their failure to satisfy CONSTANCY. In more recent work, he has given these properties a (quite prominent) role to play.²³

Another promising class of candidates is the family of properties like *being disposed to cause G experiences in some kind of observer*. Call these the *dispositional* appearance properties. These are the candidates that are endorsed in Shoemaker (2000). Like the first occurrent appearance properties, these satisfy DIFFERENCE, CORRECTNESS, and SAMENESS. Unlike the occurrent properties, they also satisfy CONSTANCY, since the dispositions are still present even when Kermit is not being observed. There's a bit of a worry about NOVELTY: these properties look a lot like the ones that dispositional theorists of color want to identify with the colors.

²³ Shoemaker (2003, forthcoming).

Unsurprisingly, this proposal still doesn't deliver CONTRARINESS; PG and PR still aren't contraries, since things can be disposed to cause G experiences in one kind of perceiver and R experiences in another. But this isn't a fatal problem. Perhaps we ought to give up on CONTRARINESS, since it seems to be incompatible with the more important requirement of CORRECTNESS. Still, if there were a proposal that would make CORRECTNESS and CONTRARINESS compatible, that would be preferable.

There is another problem with Shoemaker's dispositional candidates. Take the property PG, which is, on Shoemaker's (2000) account, the property, *being disposed to cause G experiences in some kind of perceiver*. There are two readings of this: *being disposed to cause G experiences in some possible kind of perceiver*, and *being disposed to cause G experiences in some actual kind of perceiver*.²⁴

There is trouble if Shoemaker opts for the first reading. Almost *everything* will be disposed to cause G experiences in *some* kind of perceiver or other. All it will take to be PG is to be potentially causally efficacious. But it would be nice if the property that Ernie's visual experience (and not Vert's) represents Kermit as having did more than rule out the possibility that Kermit is necessarily epiphenomenal. Kermit is fond of telling us that it's not easy being green. Surely it's not *that* much easier to be phenomenal green.

There is also a more principled worry in the neighborhood. It's not just too easy to be PG. It's also too easy to be PR. In fact, it's *exactly* as easy to be PR as it is to be

²⁴ Terminological point: it's plausible that kinds are properties, and so the *kinds* might exist necessarily. So it's best to read 'actual kind' and 'possible kind' as 'actually instantiated kind' and 'possibly instantiated kind'. But this is awkward to write, so I'll continue to write 'actual kind' and 'possible kind' in the text, with the understanding that these may be read in the more careful way described above.

PG. Not only will all and only the potentially causally efficacious things be PG, but all and only the potentially causally efficacious things will be PR, as well. This gives us the unfortunate result that, necessarily, something is PG if and only if it's PR. On the plausible view that necessarily coextensive properties are identical, it will turn out that PG *is* PR. But then the whole strategy collapses. If PG is PR, then we lose DIFFERENCE—Ernie and Vert *aren't* representing Kermit as having different properties, and so we haven't found a representational difference to ground the phenomenal difference in their visual experiences.²⁵

So Shoemaker should either disown the plausible view about property identity, or opt for the reading that says that PG is the property, *being disposed to cause G experiences in some actual kind of observer*.²⁶ This reading, too, is ambiguous. This time the ambiguity comes from two available readings of 'actual'.

On the first reading, something is PG at a world *w* iff it's disposed to cause G experiences in some kind of observer that exists in @. (Where '@' is a name for the actual world.) Call the properties we get on this reading the dispositional@ properties, and the proposal the dispositional@ proposal.

²⁵ The view that necessarily coextensive properties are identical is plausible, but it's not mandatory. Still, two points: First, a difference only with respect to which *finer*-grained properties are represented still won't deliver a difference in *pure* representational properties. Second, it doesn't seem as if we want *this* kind of consideration to motivate our choice of a theory of properties. It would be better to have an account that was neutral between theories of properties.

²⁶ Another possible response: there's a difference in which modes of presentation of PG (that is, PR) they employ. But if you're going to move to modes of presentation here, you might as well appeal to them right away, and explain the intentional difference between Ernie and Vert as a difference in which modes of presentation of *being green* they employ. (See Chalmers (2004) for this kind of proposal.) The move to appearance properties doesn't seem to have any work to do once we've got modes of presentation in the picture.

On the second disambiguation, something is PG at a world w iff it's disposed to cause G experiences in some kind of observer that exists in w . Call these properties the dispositional_w properties, and the proposal the dispositional_w proposal.

The difference between the two disambiguations is in the force of the word 'actual'. On the first reading—the dispositional_@ proposal—the force of 'actual' is to restrict the relevant kinds of observers to ones that exist in *our* world. On the second reading—the dispositional_w proposal—the force of 'actual' is to restrict the relevant kinds to ones that exist *in the world where the object is*. (There's an obvious connection to discussions of two-dimensionalism here. The dispositional_@ proposal corresponds to the *horizontal* reading of 'actual', while the dispositional_w proposal corresponds to the *diagonal* reading.²⁷)

There are two reasons why the dispositional_@ proposal cannot be right. First, even if this move *does* keep PG and PR from being necessarily coextensive, it's just by *luck*. If there are enough kinds (or just the wrong kinds) of observers in the actual world, then it will still turn out that PG and PR are necessarily coextensive. In some worlds there are enough different kinds of observers that everything that's potentially causally efficacious is disposed to cause G experiences in some kind of observer that exists there, and also to cause R experiences in some kind of observer that exists there. If the actual world is one of these, then PG and PR will be necessarily coextensive.

²⁷ See for example Davies and Humberstone (1980), Stalnaker (1978), Chalmers (1996), Jackson (1998).

In other worlds, there are only observers like Ernie and observers like Vert. If the actual world is one of these, then again PG and PR will be necessarily coextensive. An object is disposed to cause G experiences in Ernie iff it's disposed to cause R experiences in Vert, and vice versa. So if the actual world has only two kinds of observers, and they're spectrum inverted relative to each other, an object will be disposed to cause G experiences in some actual kind of observer iff it's also disposed to cause R experiences in some actual kind of observer. And so PG and PR will be necessarily coextensive, and so very plausibly identical, and so we won't have a representational difference.

We should not adopt a proposal that only works if the contingent facts about what kinds of perceivers there actually are turn out the right way. A solution to the conflict between intentionalism and the possibility of spectrum inversion ought to be *general*—whether or not it's viable ought not to hinge on these kinds of contingent facts about which kinds of perceivers there happen to be.²⁸

The second reason why this reading fails is that it delivers the wrong results for *merely counterfactual* spectrum inversion. Suppose that (as I imagine is true) the case of Ernie and Vert is counterfactual, and that there isn't any *actual* spectrum inversion—there aren't any actual observers that have R experiences when they look at green things. Then Vert is of a kind that doesn't exist in @. Kermit will be PG, because he's disposed to cause G experiences in *us* and we're (obviously) of an actually existing kind. But Kermit *won't* be PR, because while he's disposed to cause R experiences in Vert, Vert is

²⁸ Again, an appeal to modes of presentation might look promising here. See note 25 for why I think this isn't a good idea.

not of any actually existing kind. So we haven't satisfied CORRECTNESS. Vert's visual experience (in w) represents Kermit as being PR, but this is a *misrepresentation*, since Kermit's not disposed to cause R experiences in any kind of observer that exists in @. So on this reading, we can accommodate *actual* cases of spectrum inversion, but not counterfactual cases. But spectrum inversion cases don't need to be actual to demand proper treatment. We need a theory that gets the right results for both actual and merely possible cases.²⁹

What about the dispositional_w reading—the one that says that something is PG in w iff it's disposed to cause G experiences in observers of a kind that has instances in w ? A problem with this reading is that it gives the wrong results about the appearance properties of things in counterfactual situations. Take some uncooked lobster in a counterfactual world w . Call him Lenny. Lenny is green, and we're not spectrum inverted, so he's disposed to cause G experiences in us. So he's PG. Or at least, he *ought* to be PG. If something's PG in w iff it's disposed to cause G experiences in some kind of observer that exists *in* w , then Lenny *won't* be PG in worlds where there aren't any observers like us.³⁰ And this is implausible. If something's disposed to cause G

²⁹ Another problem: we also lose the plausibility of SAMENESS for me and my counterfactual inverted counterpart. It's extremely implausible that inverted-me, looking at Kermit's local (presumably still green) counterpart in some counterfactual world w and having an R experience, is representing things as being disposed to cause R experiences in some kind of observer that exists in @. That is, it's extremely implausible that his visual experience represents Kermit as having the appearance property that, according to the dispositional@ proposal, *my* experience represents ripe tomatoes as having.

³⁰ Let's suppose that lobsters' visual phenomenology, if any, is such that other (uncooked) lobsters aren't disposed to cause G experiences in them.

experiences in *us*—if, were we to look at it, we'd have a G experience—that ought to be sufficient for it's being PG. The second reading doesn't deliver this result.

This is a lot less bad than the problems for the dispositional_@ reading—if we're going to go with a dispositional story, we should definitely go with this one. It delivers a somewhat strange result about the appearance properties of merely possible things, but it's not clear how big of a problem this really is. The more serious problem is that it doesn't deliver CONTRARINESS. Similarly, while sticking with the occurrent appearance properties would avoid this whole cluster of problems for the different dispositional proposals (at the cost of giving up CONSTANCY, but maybe this is not so much of a problem), it still wouldn't get us CONTRARINESS.

We've looked at properties such as *being disposed to cause G experiences in some kind of observer* and *currently causing a G experience in some kind of observer*. These are the candidates that Shoemaker takes seriously in his discussion of appearance properties. But it's also worth looking at the prospects for some properties that Shoemaker does not discuss very much: properties where the kind of observer is fixed, such as *being disposed to cause G experiences in observers of type A*, and *currently causing a G experience in an observer of type A*. Let's call these *type-specific* properties. Type-specific properties, either occurrent or dispositional, seem like attractive candidates, because they hold out hope of delivering something like CONTRARINESS: when Ernie learns that Kermit is causing a G experience in an E-type observer, doesn't he also learn that Kermit is *not* currently causing an R experience in an E-type observer? Similarly,

when Ernie learns that Kermit is disposed to cause G experiences in E-type observers, doesn't he also learn that Kermit is not disposed to cause R experiences in E-type observers? So it looks like, though we don't get CONTRARINESS – there's no incompatibility between the property that Ernie's experience attributes to Kermit and the one that Vert's experience attributes to him – we do get an incompatibility between the property Ernie's experience attributes to Kermit and the one that Ernie's experience attributes to ripe tomatoes.

Similarly, though we don't get SAMENESS on this proposal (Ernie's visual experience attributes properties of causing or being disposed to cause certain kinds of experiences in E-type observers, while Vert's experiences attribute properties indexed to V-type observers), we *do* get some very salient *similarities* between the property that Ernie's experience attributes to Kermit and the one that Vert's attributes to ripe tomatoes. Perhaps this is enough – perhaps it was a mistake to insist on CONTRARINESS and SAMENESS, and we ought instead to be satisfied with this near-CONTRARINESS and near-SAMENESS.³¹

My first worry about this approach is that it seems a bit odd to build observer-types into the contents of our perceptual experiences – this seems like the wrong sort of thing to show up in perceptual content. But let's set this aside, and suppose that there is really no problem about building perceptual types into the contents of our experiences.

³¹ Thanks to an anonymous referee for pointing out this possibility.

The second worry is that near-SAMENESS isn't enough: securing SAMENESS is important for the success of our defense of modality-impure intentionalism. There is no obvious barrier to there being two types of observers (call them A and B) such that A-type observers visually represent the B-indexed properties. It would be weird, but not weird enough to be metaphysically impossible, for one type of observer to visually represent which things are causing (or disposed to cause) certain kinds of experiences in *other* kinds of observers. So you, as an A-type observer, might visually represent the apple as causing G experiences in B-type observers like me (even though it causes R experiences in A-type observers like you). Your experience with this content needn't be G. There is room to fight about whether this sort of case is genuinely possible, but it certainly *seems* possible, and as we will see, there is a proposal that secures SAMENESS, and not just near-SAMENESS.

Finally, and perhaps most seriously: Contrary to initial appearances, the occurrent type-specific properties don't actually deliver the advertised sort of near-CONTRARINESS, and the dispositional ones are vulnerable to the same necessary coextension problems discussed above. The occurrent properties don't secure near-CONTRARINESS because of the possibility of the following sort of case: you and I are both E-type observers, looking at Kermit. Your view of Kermit is unobstructed. Mine is obscured by a distorting haze, which makes Kermit look red to me. Kermit is causing a G experience in you and an R experience in me. You *don't* learn, when you learn that Kermit is currently causing a G experience in an E-type observer, that he isn't causing an

R experience in some *other* E-type observer who is differently situated relative to Kermit. This problem becomes more pronounced when we consider finer-grained differences in phenomenology. (The phenomenology of dark green vs. light green, of scarlet vs. crimson, etc.) (Making the observer-types extremely fine-grained and extrinsic can help to avoid this problem, at the expense of making the first worry much more pronounced, and sacrificing SAMENESS even for different experiences of the same observer over time.)

This problem doesn't arise for the type-specific dispositional properties, but a worse one does. If E-type and V-type observers are systematically inverted, then *being disposed to cause G experiences in E-type observers* and *being disposed to cause R-type experiences in V-type observers* will be necessarily coextensive, and we won't have a difference in possibility-carving content between Ernie's and Vert's experiences of Kermit. In that case, we won't have preserved our modality-impure intentionalism.

None of the candidates we've looked at so far is completely satisfactory. None of them satisfy CONTRARINESS, and many are seriously flawed in other ways. In the next two sections I will suggest that we can do better. Section 5 lays out the groundwork we'll need in order to state the proposal, and section 6 argues that my proposed candidates do better than any of the ones we've seen so far.

5. Self-Locating Content and Centering Features

I take it to be fairly well established that we don't just represent which possible world we're in, but also *where we are* in the worlds that we take to be candidates for actuality.^{32,33} We could know exactly which world is actual, and still not know whether the grocery store is to the left or to the right, or whether the car is parked nearby or far away, because we don't know *where* in that world we are. So we need to treat many propositional attitudes as attitudes not toward *possible-worlds* propositions, but toward *centered-worlds* propositions.³⁴ That is, many contentful mental states have centered worlds propositions (or things that determine centered-worlds propositions) as their contents.^{35,36}

We describe the possible-worlds content of a representation (largely) in terms of the properties that things are represented as having. By representing Kermit as having the property, *being green*, we pick out a class of worlds—all and only the worlds in which

³² See for example Lewis (1979) and Perry (1979).

³³ Read 'where we are' broadly here, so that ignorance about your *temporal* location also counts as not knowing where you are.

³⁴ A *centered world* is to a possible world what a map with a "you are here" arrow added is to an arrowless map. Centered worlds single out not just a way for the world to be, but a location within the world.

They're best thought of as ordered pairs of a world and a center. There are different ways of picking out a center—the center could be, for example, a spacetime point, or an individual, within the world. Not much hangs on this decision, but it will be convenient for present purposes to take centers to be <individual, time> pairs. Some people talk about centered worlds, others about self-attribution of properties.

Exposition is easier for centered worlds, but the same points can be made for self-attribution of properties.
³⁵ A complication: anything that determines a possible-worlds proposition determines a centered-worlds proposition. It just determines (to introduce a technical term) a *boring* centered-worlds proposition, that includes, for each world *w*, either all of the positions in *w* or none of them. So the claim is really that many contentful mental states have contents that determine *interesting* (i.e., non-boring) centered worlds propositions.

³⁶ It's worth noting that Perry *doesn't* advocate moving to centered-worlds contents for belief, but instead wants to make a distinction between different sorts of *belief-states* with the same content. So when we both believe that our respective pants are on fire, we're in the same belief-state, though we have different beliefs.

Kermit is green. When I believe that the world is a certain way, I represent some things as being green and others as being red, some things as being furry and others as being scaly, and so on. In this way I narrow down the range of worlds that I take to be candidates for actuality. (We can tell the same kind of story of other kinds of representational states—for example desires, fears, and, most relevantly, visual experiences.)

When I have beliefs not just about what the world is like, but about my location within it—when I have *self-locating* beliefs—something very similar is going on. I represent some things as being *nearby* and others as *far away*, some events as *present* and others as *past* or *future*, and some objects as being *on my foot* and others as *in my ear*. In this way I narrow down the range of possible *predicaments*—possible locations within worlds—that I take to be candidates to be the one that I am in. (Again, the same goes for other representational states with this kind of content—in particular, for visual experiences, if the content of visual experience is self-locating.)

I do this narrowing by singling out a set of centered worlds. Each centered world, remember, corresponds to a possible position within a world. Let us say that, when a particular centered world is the one that corresponds to my current position, I am *correctly located by* – in the sense that my location is fixed by – that centered world. No one is ever correctly located by more than one centered world at a time, and no one is ever correctly located by any one centered world for very long. I am, as I write this, correctly located by the centered world <@, <Egan, 5:58pm, July 28 2004>>. (Adopting

the <individual, time> method of picking out centers.) You are, as you read this, correctly located by <@,<you, whenever you're reading>>. And as you read *this*, you're correctly located by some other centered world, whose time element is very slightly later.

Attribution of properties, in the sense relevant to characterizing possible-worlds contents, is cheap. A belief, desire, fear, etc. attributes a property F to an object x iff it has a content that includes only worlds in which x is F (i.e., only worlds for which F delivers an extension that includes x). Attribution of centering features, in the sense relevant to characterizing centered-worlds contents, is equally cheap. A belief, desire, fear, etc. attributes a centering feature F to an object x iff it has a content that includes only centered worlds relative to which x is F (i.e., only centered worlds for which F delivers an extension that includes x).

The ways in which we describe centered-worlds contents are very similar to the ways in which we describe possible-worlds contents. In the possible-worlds case, one very common way to single out a set of worlds is by attributing some property like *being green* to some object like Kermit. In the centered-worlds case, a very common way to single out a set of centered worlds is by attributing some “property” like *being nearby* to some object like Kermit. At least, this looks like a tempting thing to say. But maybe we shouldn't be so quick to say that there are any such “properties” as *being nearby* or *being on my foot* for us to attribute to things.³⁷

³⁷ In talking this way, I don't intend to commit myself to any particular view about the semantics of 'nearby' in English. (And more importantly, I don't think I *do* so commit myself, any more that someone like Lewis commits himself to a view about indexical pronouns in English when discussing the self-

It's certainly true that if there *are* such "properties", they deserve the scare quotes—they're not *properties*. Properties are (or at least determine) functions from worlds to extensions. A "property" like *being nearby* won't do that. Which things are nearby—which things are in the extension of the "property", *being nearby*—depends not just on which world is actual, but also on *where you are* within the world.³⁸ That is, it depends on which centered world you are situated at. Lots of things are near me and far away from you. So *being nearby* doesn't determine a function from worlds to extensions, and so it's not a property.³⁹

This shouldn't be surprising. If it *was* a property, it wouldn't be fit to play the role that it does in determining centered-world contents, because we wouldn't be able to use it to distinguish between different positions within a given world. What it is is the *analogue* of a property for centered-worlds contents. That is, it's a function (or something that determines a function) from centered worlds to extensions, in the same way that a property is (or determines) a function from *worlds* to extensions. So given centered world *c*, *being nearby* will give us all of the things that are nearby if we're correctly located by *c*. (That is, given a world and a center, *being nearby* will give us the

locating belief that *my pants are on fire*.) All I'm hoping to do with *being nearby* is to provide an intuitively compelling example of the sort of feature that I have in mind.

³⁸ It also depends on the currently active standards of *nearness*. (Or maybe there's not just a single "property" of *being nearby*.) But let's ignore this complication for now—nothing will turn on this.

³⁹ Actually, my official view is that properties are functions from <world, time> pairs to extensions (see Egan (2004)). But that doesn't matter for present purposes; *being nearby* doesn't determine one of those, either. Why am I so sure that these things are not properties? Well, this is partly stipulative. Nothing bad happens if I allow that they're properties, but insist that there's an important distinction between two *kinds* of properties. One reason for insisting on the name, though, is that *properties* ought to be the sort of thing about which a nominalist/universalist/trope theorist debate wouldn't just be *crazy*. There's no plausibility at all to the idea that *being nearby* is some screwy kind of universal, or a set of screwy kinds of tropes.

set of things that are nearby if we're at that location in that world.) We can pick out a set of centered worlds by saying that they're the ones in which Kermit is nearby—the ones that, when plugged into the *being nearby* function, deliver an extension that includes Kermit.

We need a name for these things. I'll call them *centering features*.⁴⁰ 'Features' to indicate that they're property-analogues, not properties, and 'centering' because their role is to select not just a world, but a center (attributing a *property* to Kermit selects between worlds that might be the one that I inhabit, while attributing a *centering feature* to Kermit selects between predicaments that might be the one that I am in).

It's worth drawing attention to the fact that not all self-locating contents are *geographically* self-locating. Consider the centered-worlds propositions, *that my pants are on fire*, *that everyone is out to get me*, *that Kermit is out to get me*, and *that Kermit is disposed to cause greenish sensations in me in standard viewing conditions*. These are distinct from the corresponding possible-worlds propositions, *that Egan's pants are on fire*, *that everyone is out to get Egan*, *that Kermit is out to get Egan*, and *that Kermit is disposed to cause greenish sensations in Egan in standard viewing conditions*. These sorts of self-locating contents have nothing much to do with *geographical* self-location, and the centering features that they attribute to things (e.g., *being my pants*, *being out to get me*, *being disposed to cause greenish sensations in me in standard viewing conditions*) are importantly unlike the ones (such as *being nearby*) that I've discussed so

⁴⁰ Thanks to Alex Byrne for the name.

far in that they're not features of having a certain geographical location relative to the believer.

So just as self-locating belief isn't restricted to *geographical* self-location, centering features aren't restricted to relative-position features like *being nearby*. I'll continue to frame the discussion of centering features primarily in terms of *being nearby* and *being far away*, but this is simply because it keeps the presentation cleaner, and avoids unnecessary complications, not because centering features are one and all concerned with the geographical locations of things.

At this point, it should be pretty clear that centering features are useful things to have around for describing centered-worlds contents. In the same way that it's useful to talk about the properties that people attribute to things in describing their beliefs about which world is actual (try specifying the content of my belief that Kermit is green without appealing to properties), it's useful to talk about the centering features that people attribute to things in describing their *self-locating* beliefs (beliefs about which *predicament* they're in).

If we grant that we should take the contents of at least some kinds of representation to be sets of centered worlds rather than sets of worlds (as seems incredibly plausible), then another point becomes clear: *A difference in the attribution of centering features is a representational difference*. If my beliefs represent Oakland as *nearby* and yours represent it as *far away*, then there's a difference in the representational content of our beliefs (even though both of our beliefs might be correct). This suggests a

new set of candidates to be the appearance properties: perhaps Ernie's and Vert's visual experiences differ in which *centering features* they attribute to Kermit. I will develop this proposal in the next section.

First, though, we should pause to address some potential concerns about centering features and their place in the contents of perception. (What follows owes a large debt to the comments of an anonymous referee.)

There are, I think, two major grounds for concern about the claim that perceptual experience attributes centering features to things. The first of these is *phenomenological*: it might seem as if the phenomenology of perception tells us that everything we are aware of in perception is either a particular or a *property* of particulars. If this is so, then it seems that there is no room for centering features in the contents of perception. I think, however, that this misdescribes the phenomenology of perception. Having self-locating content goes hand in hand with attributing centering features. You can't have self-locating content without attributing centering features to things. (Since all it takes to attribute a centering feature to something is to have the right kind of self-locating content.) So if the content of perceptual experience only deals in particulars and properties of particulars, and not at all in centering features, then perceptual experience can't have self-locating content. But this seems wrong. The content of visual experience, for example, is pretty clearly self-locating. My visual experience doesn't just tell me about how things are arranged relative to each other – it also tells me something about where *I* am relative to the things that it represents. So it seems that the

phenomenology of experience, if it tells us anything about the contents of perception, tells us that the perceptual experience has self-locating content. And if it has self-locating content, then it attributes centering features to things.

The second sort of concern is *metaphysical*. To what extent is having a centering feature like having a property, and to what extent unlike? Are centering features, on account of their observer-relativity, somehow metaphysically suspect? There are a number of interesting questions in this neighborhood, that deserve more attention than I can give them here. So rather than attempting to answer these questions, I'll instead say why I think that it's okay, for present purposes, to leave them open.⁴¹

The presence of some unanswered metaphysical questions ought not to make us suspicious of centering features as contributors to the contents of perceptual experience because, as mentioned above, the attribution of centering features goes hand in hand with the representation of self-locating contents. If my visual experiences have self-locating content, then they attribute centering features to things, and vice versa. It seems clear that my visual experiences *do* have self-locating content – my visual experience tells me something, not just about the way that the world is, but about my position within it. (For example, some things are represented as being nearby and others as far away.) If visual perception *wasn't* self-locating, then vision would be a lot less useful for navigating the world than it actually is. Questions about the metaphysical status of the secondary qualities of things do not, I think, pose any threat to the claim that perceptual experience

⁴¹ I address some of these metaphysical issues more directly in Egan (forthcoming).

attributes centering features to things – we’re stuck with self-locating content, and if we’re stuck with self-locating content, then we’re stuck with the attribution of centering features.⁴²

⁴² It’s also worth contrasting centering features with some more familiar things that they bear some resemblance to. Centering features might look a bit like Fodor’s (1987) narrow contents, or Chalmers’ (2004) phenomenal modes of presentation. Both of these are narrow, character-like entities. Fodor’s narrow contents determine a *truth-condition* relative to a context, and Chalmers’ phenomenal modes of presentation determine a *represented property* relative to a context. Both make use of something finer-grained than possible worlds to determine a content – the same narrow content can determine different truth conditions, and the same phenomenal mode of presentation can serve to represent different properties, in different situations within the same world. One reason to be worried about these sorts of content-determining things is their narrowness – If we are sufficiently thoroughgoing externalist, we will be skeptical about the prospects for *any* interesting representational properties of experience being narrow.

While this might be grounds for concern about Fodorian narrow contents and Chalmersian modes of presentation (or it might not, depending on how the fights about externalism turn out), it’s not a problem for the view I’m suggesting. I make no claims either way about the narrowness or broadness of the representation of centering features. For purposes of this paper, I am completely agnostic about whether there are any narrow representational properties at all, and if so, which ones they are. So externalist concerns about narrow content are not grounds for concern about the possibility of attributing centering features to things in experience.

Still, the similarity to these sorts of views might be grounds for another sort of worry, quite independent of issues about narrow vs. broad content. If my account *is* a story of the same general kind as Fodor’s account of narrow content, or Chalmers’ account of phenomenal modes of presentation, we might be concerned about a structural similarity to the sense-data theory. (Thanks to an anonymous referee for pointing this out.) These accounts look a bit like a sense-data theory because the things that we’re immediately confronted with in experience, and which determine the phenomenal character of our experiences, are character-like things that then determine, with assistance of stuff we’re not aware of, the *content* of our experiences. This looks very much like a veil of appearances – we’re immediately confronted by things with characters, which might then isolate us in a problematic way from the contents of our experiences. That looks potentially bad.

Centering features, though, are importantly different from Chalmersian phenomenal modes of presentation, and self-locating contents are importantly different from Fodorian narrow contents. There is no nice, natural function from centering features to regular properties, and the role of centering features is *not* to determine a regular property which then gets fed into the content of the experience. (In this important respect they are unlike Chalmers’ phenomenal modes of presentation.) Similarly, there is no nice, natural function from self-locating contents to regular possible-worlds contents, and the role of a self-locating content is not to determine a regular possible-worlds content which is the *real, complete* content of the experience. In the self-locating case, unlike the Fodorian case, there’s no appeal to any intermediate content-determining thing – everything is cashed out in terms of the content. (That’s why my story is what I called in section 2 a *modality-impure* intentionalist theory. Within a modality, it’s just content doing the fixing of phenomenal character. The Fodorian and Chalmersian stories would be *MOP-impure* – what’s doing the phenomenal character fixing is not the content, but the details of how that content is getting determined.) So while the phenomenal modes of presentation and narrow contents have a second layer of

6. Appearance Properties as Centering Features

Here is a first pass at a proposal for what PG is: it's the centering feature, *being disposed to cause G experiences in me* (or possibly, *being disposed to cause G experiences in perceivers of my kind*). At least, that's the dispositional candidate. It's also easy to identify an occurrent candidate, if we prefer occurrent appearance properties (or if we just want to have both around): It's the centering feature, *presently causing a G experience in me*.

(Methodological aside: In what follows, I'll talk exclusively about the dispositional candidates, with footnotes where the parallel things to say about the occurrent candidates aren't obvious, because trying to talk about both at once makes the discussion unnecessarily convoluted. Officially, though, I'd like to leave it open whether, at the end of the day, we'll want to go with dispositional appearance properties, occurrent appearance properties, or some combination of the two, as the features representation of which grounds the phenomenal character of experience.)

It's important not to 'read through' the indexicals in the specifications of centering features like those above. It's tempting to understand the claim that Ernie's visual experience attributes the centering feature, *being disposed to cause G experiences in me* to Kermit as the claim that Ernie's visual experience attributes the property, *being disposed to cause G experiences in Ernie* to Kermit. But this would be a mistake. It's

content behind them which they might screen us off from in some bad way, the centering features and self-locating contents do not.

important that the centering feature, *being disposed to cause G experiences in me*, is something that combines with an object to determine a *centered-worlds* proposition, not a *possible-worlds* proposition.⁴³ It's also important that it's something that Ernie's visual experience represents Kermit, limes, and uncooked lobsters as having, and that Vert's visual experience represents ripe tomatoes, fire engines, and cooked lobsters as having. The *property* that we get by 'reading through' the indexical in a given case doesn't have either of these features.⁴⁴

This meets all of the desiderata and avoids all of the problems of the other accounts. Let's look at each desideratum in turn:

DIFFERENCE: If PG is the centering feature, *being disposed to cause G experiences in me*, and PR is *being disposed to cause R experiences in me*, then PG is plainly not the same centering feature as PR. The set of possible predicaments in which Kermit is disposed to cause G experiences in me is clearly not the same as the set in which he is disposed to cause R experiences in me. So DIFFERENCE is satisfied. (The parallel for the occurrent appearance properties should be obvious.)

CORRECTNESS: Neither Ernie nor Vert is misrepresenting Kermit—Kermit really *is* disposed to cause G experiences in Ernie, and he really is disposed to cause R experiences in Vert—so CORRECTNESS is satisfied. (Again, it's clear that this also holds for the occurrent candidates.)

⁴³ That is, it's important that it combine with an object to determine an *interesting* centered-worlds proposition, to go back to the terminology of note n again.

⁴⁴ As a consequence, these properties are also bad candidates to be the appearance properties. Their prominent defect is that they don't satisfy SAMENESS.

NOVELTY: If we agree that Ernie's and Vert's experiences agree on their color content, then these centering features can't be the colors. Also, it seems quite plausible (though it's not universally agreed to) that the colors of things are perfectly objective properties of them, and that which colors things have is not observer-relative.⁴⁵

(If we thought that the colors of things *were* observer-relative, then we shouldn't have been worried in the first place—we should have resisted the claim that it was possible to have spectrum inverted subjects whose visual experiences agreed on the colors of things.)

SAMENESS: The centering feature that Ernie's visual experience attributes to Kermit (and to unripe tomatoes) is the same one that Vert's attributes to, e.g., cooked lobsters and ripe tomatoes.

CONTRARINESS: When Ernie learns that Kermit is PG, he learns that he's not PR (just as when he learns that San Francisco is *far away*, he learns that it's not *nearby*). So CONTRARINESS is satisfied, as well.⁴⁶

This is compatible with *Vert*, when he learns that Kermit is PR, learning that he's not PG (just as Vert could, if he lived in Oakland, learn that San Francisco is not *far away* by learning that it's *nearby*). In the case of centering features (like *nearby* and *far away*),

⁴⁵ In the case of the occurrent appearance properties, it's even clearer that NOVELTY is satisfied. Though this was equally clear for Shoemaker's occurrent candidates.

⁴⁶ We actually need to add an 'in present circumstances' qualification to the specification of PG in order to guarantee CONTRARINESS. If we read *being disposed to cause G experiences in me* as meaning, *being disposed to cause G experiences in me in some circumstances or other*, then PG and PG are compatible—the same thing could be disposed to cause G experiences in me in one sort of circumstances, and R experiences in some other sort of circumstances. So officially, PG should be: *being disposed to cause G experiences in me in present circumstances*. Thanks to John O'Dea for pointing this out.

unlike in the case of properties (like *green* and *red*), CONTRARINESS and CORRECTNESS are compatible.

CONSTANCY: They're dispositions, so there's no problem about Kermit's having them even when he's not being observed.⁴⁷

I conclude that centering features are the best candidates to play the appearance-property role. If we are going to go in for a Shoemaker-style solution to the conflict between intentionalism and the possibility of spectrum inversion, we should say that the representational difference that grounds the phenomenal difference between inverted subjects is a difference with respect to which centering features (of the kind described above) are represented by their visual experiences.

Unfortunately there is a lingering problem about deviant dispositions (pointed out by Alex Byrne). Suppose that Kermit was, in addition to having his usual effects on Ernie's perceptual system, was also disposed to cause Ernie to hallucinate a ripe tomato hovering in the air three feet to Kermit's left. This would be strange, but it's certainly not impossible. In this case, two bad things would happen. First, Ernie would be correct in representing Kermit as PR, since Kermit really would be disposed to cause R experiences in Ernie. This is bad, because only red things ought to be PR to non-inverted observers. Second, Ernie would be correct in representing Kermit as *both* PR and PG, (since he

⁴⁷ This is, of course, a place where the dispositional and occurrent candidates come apart.

really is disposed to cause both G and R experiences in Ernie), which shows that the first pass proposal doesn't really give us CONTRARINESS.⁴⁸

Two strategies for revising the first pass proposal to avoid this problem suggest themselves. First, we could say that something is PG iff it's disposed to cause G experiences in me in a *non-deviant way*, and start looking for some halfway satisfying way of cashing out 'non-deviant'. This will probably be hard. Second, we could say something quite a bit fancier about what it is that PG things are supposed to do—if we're fans of visual fields, the thing to say is something like, *an object is PG iff it's disposed to cause the bit of the visual field responsible for representing it to go green*'. I suspect that it is possible to make good on this *sort* of proposal without actually committing oneself to green' patches on visual fields (by talking about very specific phenomenal properties of experiences and the relations of similarity between them), but I won't argue for that here.

To sum up: Centering features do a better job of playing the appearance property role than do any of Shoemaker's candidates, because they can meet all of the desiderata. Their principal advantage is that they offer a way of reconciling CORRECTNESS and CONTRARINESS. There is an unfortunate complication about non-standard ways of causing the relevant sorts of experience, which will take a bit of work to fix. This is, however, a problem that's common to pretty much every dispositional or causal theory of

⁴⁸ An extremely similar objection also goes through against my occurrent candidates. The only difference is that, rather than appealing to deviant dispositions, the objection appeals to present-tense deviant causal chains. (Kermit's causing, right now, the usual kind of G experience, as well as a hallucination of a ripe tomato three feet to his left. So he's causing both a G experience and an R experience in me now, and so it looks like I have to say that he's both PG and PR.)

anything, and it does seem to be tractable (though it might turn out that tractating it brings along some unwelcome theoretical commitments).

Conclusion

Appearance properties were introduced to play a certain role. They're supposed to mark a difference in how spectrum inverted subjects represent colored objects like Kermit. I've argued that the best candidates to play this role aren't *properties*, but *centering features*. So while it's not clear whether or not there are any appearance *properties*—since there are things that play the role, but they aren't properties—the intentionalist should be happy. There *is* a representational difference between Ernie and Vert: their visual experiences attribute different centering features to Kermit. And as far as reconciling intentionalism with the possibility of spectrum inversion without misrepresentation goes, it's just the *presence* of the representational difference, and not the details of what *kind* of difference it is, that matters.

References:

- Boghossian, P. and Velleman, J.D. (1989) Colour as a Secondary Quality, *Mind* **98**: 81-103.
- Byrne, A. and Hilbert, D. (1997) Colors and Reflectances, in Byrne, A. and Hilbert, D. (eds.) *Readings on Color, vol. 1: Philosophy of Color*, MIT Press, Cambridge, Mass.
- Byrne, A. (2001) Intentionalism Defended, *Philosophical Review* **110**: 199-240.
- Chalmers, D. (1996) *The Conscious Mind*, Oxford University Press, Oxford.
- Chalmers, D. (2004) The Representational Character of Experience, in B. Leiter, ed., *The Future for Philosophy*, Oxford University Press, Oxford.
- Davies, M. and Humberstone, L. (1980) Two Notions of Necessity, *Philosophical Studies* **38**:1-30.
- Egan, A. (2004) Second-Order Predication and the Metaphysics of Properties, *Australasian Journal of Philosophy* **82**: 48-66.
- Egan, A. (forthcoming) Secondary Qualities and Self-Location.
- Fodor, J. (1987). *Psychosemantics*, MIT Press, Cambridge Mass.
- Harman, G. (1990). The Intrinsic Quality of Experience, *Philosophical Perspectives* **4**: 31-52.
- Hilbert, D. and Kalderon, M. (2000) Color and the Inverted Spectrum, in *Color Perception, Vancouver Studies in Cognitive Science* **9**.
- Jackson, F. (1998) *From Metaphysics to Ethics*, Oxford University Press, Oxford.
- John, J. (2004) *Presence to Mind: Phenomenal Properties and the Question of Direct Realism*, PhD dissertation, MIT.
- Lewis, D. (1979) Attitudes *De Dicto* and *De Se*, *Philosophical Review* **88**:513-43.
- Perry, J. (1979) The Problem of the Essential Indexical, *Nous* **13**:3-21.
- Shoemaker, S. (1982) The Inverted Spectrum, *Journal of Philosophy* **79**: 357-81.
- Shoemaker, S. (1994) Phenomenal Character, *Nous* **28**:21-38.
- Shoemaker, S. (2000) Phenomenal Character Revisited, *Philosophy and Phenomenological Research* **60**:465-7.
- Shoemaker, S. (2001) Introspection and Phenomenal Character, *Philosophical Topics*
- Shoemaker, S. (2003) Content, Character, and Color, *Philosophical Issues* **13**:253-278.
- Shoemaker, S. (forthcoming) On the Ways Things Appear
- Stalnaker, R. (1978) Assertion, *Syntax and Semantics* **9**: 315-332. Reprinted in Stalnaker (1999). Page references are to this reprinting.
- Stalnaker, R. (1999) *Context and Content*, Clarendon, Oxford.
- Stalnaker, R. (2000) Comparing Qualia across Persons, *Philosophical Topics* **26**: 385-405.
- Thau, M. (2001) *Consciousness and Cognition*, Oxford University Press, Oxford.
- Yablo, S. (1995) Singling Out Properties, *Philosophical Perspectives* **9**: 477-502.