

# Perception from the First-Person Perspective

Robert J. Howell

*Abstract:* This paper develops a view of the content of perceptual states that reflects the cognitive significance those states have for the subject. Perhaps the most important datum for such a theory is the intuition that experiences are 'transparent', an intuition promoted by philosophers as diverse as Sartre and Dretske. This paper distinguishes several different transparency theses, and considers which ones are truly supported by the phenomenological data. It is argued that the only thesis supported by the data is much weaker than those typically considered in the literature, and has no obvious implications for the existence of sensory intermediaries. It does, however, have implications for the content of perceptual experience. It is argued that combining the first personal project with the transparency intuition yields an error theory of perception. The counterintuitive implications of this error theory are mitigated, however, by an account of perceptual beliefs which allows for them to be true even if the perceptual states are non-veridical.

## 1.

I'm sick of working and go on a stroll. Disparate universes draw me out of my head and into the world. I pass the West Indian Bakery redolent of Jerk Chicken, the sign flashes a red hand then a white walker, and bass pumps from the cars waiting for me to cross. In the park the gray of the city is replaced by twinkling greens. Eventually the sun tires and the air washes over with a chill. The grass darkens as the shade from the trees stretches over the fields.

It's a little paradoxical that only by leaving my quiet study and entering a realm of noises, colors, tastes, and smells can I clear my head. But that's certainly the way it seems. What is doubly peculiar is that at least on a typical view of the mind I am getting out of my head by populating it with a gamut of sensations: sensations of cold, of sound, and of color.

But it certainly doesn't appear that way. The unreflective path through the world doesn't include sensations or mental entities at all—at least not dressed as such. It includes the greenness on the trees and the redness of the street sign. These things are not mine, they are not 'of' me: I am a mere witness to them. If it is true at all to say that I am having sensations, I look right through them.

## 2.

These sorts of observations are taken to support many different positions in the philosophy of perception, often without any further argument. In the name of 'transparency', 'diaphanousness', or 'directness', positions are defended that bear on the metaphysical status of perceptual and sensory states, and the

epistemology and content of perception. It is my view that the importance of this phenomenological data is often exaggerated but that for a certain sort of 'first-personal' project in the philosophy of perception it can be quite important. By a first-personal project, I mean a project that respects the authority of the subject's perspective, in at least two ways. First, it should, where appropriate, pay heed to the phenomenology of perception. If some form of 'transparency thesis' about perception is true, for example, this should be reflected in the account of the nature of perception. Second, and relatedly, the account should do justice to the 'cognitive significance' of perception to the subject. It would be in tension with this project, for example, to ascribe perceptual content to a subject that conflicts with his or her implicit understanding of their content. If an articulation of a subject's perceptual content would generate a good deal of surprise to that subject, for example, that is a sign that the cognitive significance of the perception for the subject is not being respected. A 'first-personal' theory of perception is not mandatory, of course. It might be, in fact, that in the end we have some theoretical reasons for adopting a more 'third-personal' point of view both in a scientific study of perception and in attributions of perceptual content. Nevertheless, the first-personal project provides what many of us look for in a philosophy of perception, and I suspect that its motivations are at work, sometimes behind the scenes, in the current literature. As such, it is important to see where this project takes us, and if it takes us to a different place than the more 'third-personal' project, it will be philosophically valuable to see where we can achieve a mesh between the resulting views.

In this paper I will present the view of perception that seems to be recommended by the first-personal desiderata. I will begin in section 3, by considering various claims stemming from the transparency intuition, exploring their ambiguities, and arguing that most of them are not supported by the phenomenological data. I will, however, endorse a limited transparency claim that has fewer implications. In sections 4 and 5, I will argue that a 'first-personal' approach has important implications for the content of perception. In particular, I will argue, a sort of projectivism seems to be the natural view if both transparency and the first-personal program are accepted. Projectivism and error theories are, however, rather unpopular and seemingly quite infelicitous. In section 6, therefore, I will look at a few of the reasons that have been given for rejecting error theories. I will argue that these reasons don't ultimately tell against a theory of perceptual content so long as (a) there is an easy way to recover our third-personal judgments of veridicality, and (b) the systematic errors in perceptual content do not generate systematic error in perceptual belief. I will end by showing that both of these conditions can be met, largely by distinguishing between perceptual content, perceptual aptness, and perceptual belief.

### 3.

Experience is transparent, at least in some sense. The important question, though, is whether interesting philosophical conclusions can be drawn from this

phenomenological observation. Historically, a wide variety of conclusions have been drawn. Jean Paul Sartre claimed, among other things, that the ego is a transcendent entity and not within consciousness at all, because 'if it existed . . . it would slide into every consciousness like an opaque blade' (Sartre 1957: 40). Purging consciousness of the ego, meanwhile, 'recovers its primary transparency. In a sense, it is a nothing, since all physical, psycho-physical, and psychic objects, all truths, all values are outside it; since my me has itself ceased to be any part of it.' (Sartre 1957: 93)<sup>1</sup> While Sartre focused on pure, non-reflective consciousness, many Anglo-American philosophers have emphasized that even in reflection, or introspection, the stuff of the mind is too slippery to grasp. G. E. Moore famously held:

Though philosophers have recognized that something distinct is meant by consciousness, they have never yet had a clear conception of what that something is . . . The moment we try to fix our attention upon consciousness and to see what, distinctly, it is, it seems to vanish: it seems as if we had before us a mere emptiness. When we try to introspect the sensation of blue, all we can see is the blue: the other element is as if it were diaphanous. (Moore 1903: 450)

Similarly, much later we find Gilbert Harman:

When Eloise sees a tree before her, the colors she experiences are all experienced as features of the tree and its surroundings. None of them are experienced as intrinsic features of her experience. Nor does she experience any features of anything as intrinsic features of her experience. And that is true of you too. (Harman 1990: 31)

And Harman has been followed by many others, including (Tye 2000, 113: 51–2) and (Dretske 1995: 62), who also insist that experience is transparent.

What to make of these observations? They reflect undeniable phenomenological data, but there are many different claims that could be made in light of that data. Here are several possible theses:

*Unreflective Phenomenological Transparency:* when we are not reflecting on our mental states, but are engaging with the world, we are not aware of our sensory states but of the qualities and things in the world.

Sartre dwells on this unreflective transparency.

Harman and others, however, do not confine themselves to talking about unreflective episodes. They seem to argue that even when we reflect, or attempt to introspect our mental states, we are not directly aware of them but are constantly led back to the qualities and objects that the states present. At most, on this view, we are aware of our experiences indirectly, in virtue of our awareness of the objects in the world that they 'present'. This suggests the following stronger thesis:

*Reflective Phenomenological Transparency:* one is never directly aware of one's own experiences.<sup>2</sup>

Perhaps the most important and most overlooked ambiguity in these theses arises because it is possible to be aware of something, and even attend to it, without being aware of it as the thing that it is. When waking from a nightmare, I've seen my bathrobe hanging on the back of the door and to my tired eyes it looks as though a menacing stranger is hovering in wait. I can stare right at the robe, without being aware of it as a robe. It is conceivable that we are aware of our sensations in the same way, and are in fact attending to them, only without being aware of them as sensations. It might even be terribly difficult to become aware of them as sensations, due to the exigency of the things in the world that they present, but it still might nevertheless be accurate to say that we are aware of them.

These considerations suggest a further disambiguation of the transparency thesis that cuts across both reflective and non-reflective transparency:

*Strong Transparency:* We are not directly aware of our experiential states, not only as experiential states, but at all.

*Naive Transparency:* We are not directly aware of our conscious experiential states as experiential states.

This, then, gives us four different transparency theses: Strong Unreflective Transparency, Strong Reflective Transparency, Naive Unreflective Transparency, and Naive Reflective Transparency.

These transparency theses, though, suffer from another vagueness. What constitutes direct awareness? Following Jackson and others, we can say that one is directly aware of an object O when one is aware of it, but not in virtue of being aware of some distinct object S (Jackson 1977: 19–20; Stoljar 2004). But what does it take to be aware of something?

On one very narrow notion of 'awareness', one is only aware of an object if it is the object of a perceptual state. If this is the notion at play in the transparency theses, they would be upheld so long as what we call 'introspection' is not really a perceptual relation, or, even weaker, if in the process of perceiving the world it is not correct to say that we are perceiving our experiences.<sup>3</sup> This defense of the transparency theses is available, but it threatens to make transparency a terminological issue.<sup>4</sup> If the transparency theses claim that we are not aware of our experiences or sensations in such a technical sense, it doesn't seem as if the much touted phenomenological observations are doing any work.

Phenomenology might, however, be taken to comment on a more colloquial, common-sense notion of awareness. In this broader sense, one is aware of something simply if it makes a cognitive impact by being present to one's conscious mind. One way something can make such a cognitive impact is by being perceived in some technical sense, but that isn't the only way things are present to the mind. If I am wearing rose-colored contact lenses, I am visually aware of the color of the lenses. I do not perceive the color, at least in a more restrictive Shoemakerian sense, which requires that I have perception of an object and that this perception provides me identification information for that object.<sup>5</sup>

It might even get to the point that I cannot easily attend to the color or recognize I am aware of it, but it is making an important cognitive—and visual—impact nonetheless. The same cannot be said of the air between my eyes and the computer screen, or even my clear contact lenses. The phenomenological data speak to this broader sense of awareness. The transparency theses are best construed as maintaining that we cannot be aware of our sensations and experiential states in this sense, not in a more technical sense.

The phenomenological intuitions really only support the naive versions of the theses. The data leave open the possibility that we are in fact directly aware of our experiences, but that we do not realize that they are experiences. Much as I can be aware of my robe without recognizing it as a robe, we could be aware of our experiences but not as modifications of ourselves. As the projectivist view I later propose will hold, our experiences might appear to us as the world's veneer, and not as modifications of our own minds. There might be other arguments against that view, but the phenomenology does not speak against it. It would therefore be going beyond the phenomenological evidence to assert the strong versions of either the reflective or the non-reflective transparency theses.

Of the naive transparency theses, the unreflective thesis is most easily supported by the data. In an unreflective mode, we are rarely aware our sensations and experiential states as mental states. Sartre seems correct on this front: at least in the mode of engagement with the world, our mind doesn't make appearance as a mind at all, but is merely the recipient of presentations. In reflection, or when perceptual anomalies focus our attention on how things are subjectively, the situation seems different. If one drinks orange juice after brushing one's teeth, the resulting sensations can be disgustingly potent. One might say that the orange juice tastes awful, so even here one's impressions can point outward, but there is an edge to that sourness that does not present itself as a modification of the juice. Much as the warmth of a candle flame shades into the pain on one's palm, so many sensations that initially point outward eventually make their subjectivity apparent. If one watches a blue television screen in the dark for some time, and then turns on the lights to look at a white wall, there is a very clear afterimage. (Block 1990, 2005).

Many philosophers have presented examples such as these to counter transparency claims, and the defenders of transparency, perhaps most notably Michael Tye, have been rather ingenious in responding.<sup>6</sup> One can easily imagine Tye arguing that the sourness in the orange juice involves a veridical representation of a chemical reaction on one's tongue, while afterimages are non-veridical representations of a portion of blue on the wall.<sup>7</sup> But it is important to be clear about what these responses really show. Typically, these responses involve pointing out that the adduced cases are not really counterexamples to transparency because the experiences can still be described in terms of the way the world is represented as being. That this is the case, however, does not mean that we are not aware of these intermediaries, and aware of them as sensations. When having a blue afterimage, it might in fact be possible to describe the experience in a representational idiom, such as 'the wall has a blurry blue spot' or even 'a

sheer blue cloud is in the air between the wall and me', but the fact that this can be done does not mean that I am not aware of these sensations, and even directly aware of them. It might be, in fact that we are directly aware of such sensations, and that in the normal instance they refuse to show their true faces but instead appear simply as modifications of the world.

Strong transparency claims that we become aware of our experiences, if at all, only derivatively by focusing upon the objects of perception. Something seems right about this. It's important to notice, however, that this claim has two parts. One part maintains that in every case where one has knowledge of one's own experiences, one must also be aware of the object of those experiences. Call this the world/experience thesis. This thesis is not enough to establish that one is not directly aware of one's experiences, because this thesis says nothing about directness. There are (at least) two different pictures which fit the world/experience thesis, but only one of them supports a strong sense of transparency. On the first picture one is not aware, in any sense, of experiences, but one can, based on an inference, conclude based upon the appearance of outer objects that one is having such and such an experience. Call this the world-first model. On this model, introspection is really perception plus an inference.<sup>8</sup> Consider an analogy. When my eyes were younger, I only had trouble seeing when driving at night. As the day went on I might forget whether I had inserted my lenses in the morning, and since they were transparent and comfortable I was not aware of them. One way I could discover I was wearing them was to get on the road at dusk: if the road signs were legible, then I could infer I was wearing my lenses. Something like the world-first model seems to be operative here.

This isn't the only model that fits the world/experience thesis, however. It is also possible that one cannot be aware of one's experiences without being aware of the experiential objects because experiences ineluctably present the objects, but that nevertheless it is only by being aware of those sensations in the first place—though not perhaps as sensations—that the objects make their appearance. This is the experience-first model. A better analogy here, it seems, is the case of an excellent television set.<sup>9</sup> Suppose the television set is very large, such that its screen filled one's field of vision.<sup>10</sup> Suppose the set is broadcasting a picture of a tree. I am surely not directly aware of the tree. I am only aware of the tree by being aware of the pixels of the set—even if I'm not aware of them as pixels. Now it might in fact be that I cannot look at the TV in its current state without being aware of a tree. Perhaps studying the details of the tree is the best or in fact the only way to ascertain the quality and resolution of the set. Nevertheless, my perception of the tree is mediated, and it is mediated by something of which I am more directly aware (Brown 2009). The world-first model doesn't hold here, and neither does strong transparency. The world/experience intuition is silent about which model to accept, and since the phenomenological data at best support the world/experience intuition, the phenomenological data do not force us to accept strong transparency.<sup>11</sup>

The upshot of this discussion is that the phenomenological data leave it open that we are, in fact, directly aware of our experiences even in unreflective cases,

because it leaves it open that we are aware of them but not as experiences. What of a weaker transparency thesis that we cannot be directly aware of our experiences as experiences, or our sensations as sensations? This view does not fare well either. In cases such as drinking orange juice after brushing one's teeth, seeing blurry images and afterimages, and having ringing in one's ears, it seems clear that sensational intermediaries become apparent. Even if there are ways to describe these experiences representationally, and even if a naive perceiver might not recognize them as experiences, this is no argument against the claims that the anomalous situations do, in fact, lead many perceivers to be aware of their sensations as sensations, and in becoming so aware the perceiver need not be aware of those sensations in virtue of being aware of anything else.

The naive transparency thesis, however, seems correct. It is clear that in an unreflective, naive mode, we are not aware of our sensations as sensations. Naive Transparency is not in conflict with the thesis that we are, as a matter of fact, aware of our sensations in the normal, unreflective course of perception. It is only in conflict with the idea that we are aware of them as intermediaries. I think this gets at the kernel of truth in the transparency theses. This suggests an interesting and puzzling feature of experience. If the experience-first picture is correct, and sensations are intrinsic features of our minds, it is somewhat peculiar that they refer us ineluctably to objects. Without drawing attention to their natures, they usher attention to the world. In the unreflective mode sensations are not cognitively salient—whatever role they play, they play it without their role as sensations being cognitively significant. While this has little bearing on the *metaphysics* of perception, it might have significant bearing on the *content* of perception.

#### 4.

When we have conscious experiences, there is something it is like to have them: a certain feeling or subjective impression. The experience associated with seeing a red apple feels different from the experience of seeing a green apple, and the experience of heat feels different from the experience of cold. Call this feeling the *phenomenal character*<sup>12</sup> of the experience. As we've noted, most if not all experiences present the world in a certain way. Not only does it feel a certain way to have a heat sensation, but that sensation represents the world as being a certain way, and the perceptual experience can be correct or incorrect in its representation. Call this feature of experiences the *representational content* of experiences. An obvious question now arises: what is the relationship between the representational content of an experience and the phenomenal character of the experience?

On one extreme view, the answer would be that what an experience is like for the possessor makes no contribution to what the experience represents. Just as it doesn't matter to the representational properties of a photograph if it is printed on paper or displayed on an LED display, so the phenomenology of a mental representation does not affect its representational properties. Let's say that such

a view upholds *perceptual content* but denies *phenomenal content*. A more moderate view would say, on the other hand, that there is a connection between representational content and phenomenal character, but that the relationship is contingent. The fact that a certain state feels a certain way for a subject does not by itself necessitate that the subject is representing anything, but given certain connections between states of that type and the world, a subject represents something in virtue of having that state. On this view, phenomenal characters are like the words in a written language. A particular word, such as 'dog' represents something, but not simply because of the shape of the letters or characters that make up 'dog'. 'Dog' represents dog because of contingent features of 'dog' type markings—in this case, contingent linguistic conventions. Although it departs from some usage in the literature, it makes sense to say that such a view accepts a sort of *phenomenal content*.<sup>13</sup> What it does not accept, however, is that there is a type of representational content that supervenes upon, or is necessitated by a state's having a particular phenomenal character. Let's say that this stronger view maintains that there is *intrinsic phenomenal content*.<sup>14</sup>

Which view should we accept? It is difficult to argue conclusively here, so I will argue conditionally. To the extent that one takes the transparency theses to have any bearing upon the issue of experiential content, one should embrace phenomenal content of some sort. To let one's theorizing about content be guided by such considerations, though, is to embrace a certain deep connection between the way that things seem phenomenally for the subject, and the way the subject represents things. But once one embraces this, it seems natural to say that not only is there phenomenal content, there is intrinsic phenomenal content.

As I walk across the campus in spring, there is something it is like for me to see the grass—I naturally say it seems greenish. I don't, of course, think of the green as being anywhere but in the grass, but the greenishness is part of the phenomenal character of my experience. There is something different it is like to look at the portion of the lawn worn down by the students. This experience has a brownish character. Again, this is part of the phenomenal character of my experience, but as the naive transparency thesis tells us, it is not presented as such. I look through it, straight to the world. Because of these differences in character, I see the grass to be different in one place rather than another. In light of this, to deny the existence of phenomenal content would be to say that the greenish and the brownish characters of my experiences do not actually present the world as being a certain way; that the world does not have to be a certain way for that to be the correct way for my experience to seem. From a certain detached perspective, which sees phenomenal consciousness just as another spoke in our cognitive wheels, perhaps this is plausible. But it is surely not plausible from my point of view as the conscious perceiver. The phenomenal quality of my experiences provides my access to the world. This is the most important difference between me and a blind-sighter: the world seems to be a certain way to me because my experiences are a certain way. Suppose my visual field suddenly went black, but I retained all of the old informational connections. I might or might not withhold judgments about the way the world is. Never-



theless, it would seem to me as though a layer of my representational apparatus had been removed. It might be that some non-phenomenal representations are still there, but my reaction would not be just that I seemed to feel different. I would think that the world no longer seemed to be, at least visually, any way at all. The world seems to be a certain way for a subject at least partly because of the phenomenal states the subject is in.<sup>15</sup>

An experience's seeming a certain way phenomenally is connected with the world's seeming to be a certain way.<sup>16</sup> The question now is whether this is a contingent connection. At least given the motivations we have been pushing, it's difficult to see how the connection could be contingent. Once one accepts the phenomenological motivation for phenomenal content, one accepts that there is a very close connection between how things are phenomenally for the subject, and how the world seems to be to the subject. Whatever contingent connections hold between a phenomenal state and the world are behind the scenes, as far as the subject is concerned. The presentational aspect of the phenomenal is, however, available.

These intuitive considerations can be fleshed out by a few arguments. Suppose two items, a and b, both produce the same exact phenomenal character in me—say, GREEN the phenomenal state I get when I look at green grass. Now suppose a does so in virtue of property F, and b does so in virtue of property G, and F and G are not two determinates of a single property—they are wholly distinct natural kinds, say. Let's also suppose that GREEN is connected to both G and F by whatever contingent connection C the contingent phenomenal content theorist thinks must be added to phenomenal character to get representation. The question is, will it seem to me that there is something in common between a and b? It seems clear that it will. But there is nothing in common between a and b. So do I misrepresent something? It seems that this theorist must say no, since the 'representation-making' contingent connections hold between GREEN and both G and F are both present. This is untenable.<sup>17</sup>

Consider another case. Phenomenal characters come in determinates and determinables. A SCARLET phenomenal character and a VERMILLION phenomenal character are both determinates of the determinable character RED. Suppose that SCARLET bears C to F, and VERMILLION bears C to G, and F and G are not both determinates of a single determinable property. Now suppose I look at a that is F and b that is G. It will seem to me that a and b have something in common. This is a misrepresentation. But the contingent phenomenal content theorist cannot account for this, since there is no property with which RED bears C. The representation of the determinable property stems from the phenomenal characters themselves, regardless of the existence of a further relation.

The general intuition here is that the relation between phenomenal character and phenomenal content is closer than the connection between, for example, a word and what it represents. We can easily swallow the notion that when I write 'rot' and a German writes 'rot' we are representing completely different properties. But can we accept that when I have an experience with a phenomenal character of GREEN and someone else has a phenomenal character of GREEN

that we are not representing the world as being the same way? Perhaps, with a stretch, but only by abandoning the first-personal point of view that made phenomenal content seem so plausible in the first place; only by closeting the intuition that our experiences seem to refer us directly to the world.

## 5.

I've argued that that the reasons that lead us to believe there is phenomenal content at all provide us with reasons to believe in intrinsic phenomenal content. But now what is represented by these phenomenal characters? What is the relevant phenomenal content? The particular contents will of course depend on the particular characters. Following in a long tradition, let's focus on visual experience. It seems conceivable, and perhaps even likely, that not all perceivers have the same subjective experiences as a result of the same external stimuli. It seems possible, for example, that due to internal differences—in the eyes or brain, perhaps—Jack and Jill actually see things very differently. When Jack sees a red apple, he has the sort of sensation that Jill has when she sees a green apple, and when Jill sees a red apple she has the sort of sensation that Jack has when he sees a green apple. Jack and Jill are red/green inverters. Despite the fact that things are quite different for Jack and Jill they will by and large be behaviorally indistinguishable. They will both call tomatoes 'red' and grass 'green'. Jack and Jill both go when the green light is lit, and stop at the red light. The inversion only makes a difference to the subjects from the inside. Given this, the question arises as to what should be said about the correctness of Jack and Jill's perceptions. Either they both perceive color veridically, or only one of them does, or neither of them does.

There seems little basis to say that one of them does while the other does not. It is, perhaps, tempting to say that whichever one represents things in the way that we do—we normal perceivers—is the one perceiving correctly. This, of course, raises the question of what makes one set of perceivers normal. There are, of course, various possible answers—that the majority of perceivers constitutes the norm, that normal perceivers are those that perceive in accordance with evolutionary design, etc. But these suggestions seem largely beside the point. Given that they behave similarly, and both get along perfectly well in their environment, it seems doubtful that there is a survival benefit that would accrue to a Jackish perceiver over a Jillish perceiver or vice versa. Similarly, the suggestion of the average perceiver just pushes the question back a bit—one can just as easily imagine that there is no majority among perceivers: half the population could be Jackish and half could be Jillish. So, the question remains, which perceives correctly?

There simply seems no principled basis for deciding that one of the inverters is correct and the other is not.<sup>18</sup> So, either they are both correct, or neither is. There is an intuitive argument for denying that they both perceive correctly. Either they are both representing the same property, or they aren't. The idea that they are

both representing the same property is in tension with the very intuition behind phenomenal content. It was intuitive to say that there was phenomenal content precisely because veridicality conditions seem to stem from the way the world phenomenally appears to a subject, and the world phenomenally appears differently to Jack and to Jill. So it seems they must be representing different properties. But if they are both perceiving correctly, and assuming it is colors they are representing, when they are looking at the very same spot on a tomato there must be two different colors in one place at the same time. This seems counterintuitive enough, but one can't keep the number of colors in one place down to a mere two. It seems, in fact, that there will wind up being a color for each type of imaginable perceiver. We can imagine, in addition to Jack and Jill, Jim who sees things like Jack but a shade or two darker, and Jane, who perceives things like Jill, but a shade or two lighter, etc. On the face of it, this seems to be multiplying colors beyond necessity. The argument is thus:

1. There is no principled reason for saying that Jack's perception is veridical while Jane's isn't, or vice versa.
2. They cannot both be veridical, since objects cannot have two fully determinate colors in the same place.  
Therefore, neither Jack nor Jill's perception is veridical.

The result of this argument, of course, is that none of our perceptions of color is veridical, since there is no difference in principle between us and Jack and/or Jill.

An obvious response to this argument is to maintain that it only holds for a naive view of colors which, given the possibility of inversion, should be rejected. Instead, we should adopt a view according to which colors are dispositions to produce phenomenal appearances in perceivers.<sup>19</sup> Thus, premise 2 can be rejected and objects can have two fully determinate colors in a single place. At a particular place the apple can be red-for-Jill and red-for-Jack. Once the subjective indexing of colors is made explicit, the idea that two objects can have two determinate colors in one place is not so absurd.

While this move is tempting, it seems to conflict with some of the intuitions behind phenomenal content (Chalmers 2006: 56). The point can be made by pairing the possibility of spectral inversion with Ned Block's Inverted Earth (Block 1996). It seems quite possible that there be a world where the grass reflects light in the same way red apples do on earth, and blood, hearts, and stoplights reflect light the way that grass does on earth. It also seems possible that JackT and Jill are inverted, and that Jill lives on earth and JackT lives on inverted earth. In this case when JackT looks at blood and Jill looks at blood, they will see the blood in the same way—that is, it will seem the same to them. Blood looks to JackT on inverted earth, just like blood looks to Jill on regular earth.

If there is a type of representational content that is fixed by phenomenal character, or what it is like for the subject, then the way blood represented as

being by JackT is the same as the way things are represented as being by Jill. This Twin-Argument is thus:

- 1\*. JackT and Jill either both perceive veridically or neither does.
  - 2\*. The color of blood on Inverted Earth seems the same to JackT as the color of blood on Earth seems to Jill.
  - 3\*. By hypothesis, the color of blood on Inverted Earth is not the same as the color of blood on Earth.
- Therefore, neither JackT nor Jill is perceiving veridically.

This, again, suggests that we are in the same boat as JackT and Jill since there is no principled difference between us and them.

Adopting dispositionalism about colors doesn't help with this case since on that view JackT and Jill would be seeing different properties. This would preserve veridicality, but only at the cost of intrinsic phenomenal content—it would no longer be the case that there is a content that supervened upon phenomenal character. Despite the fact that the colors seem the same phenomenally to the two of them, Jill is not really representing green things in the way that JackT is representing red things. Jill is representing blood as (that which causes RED in Jill) and JackT is representing twin-blood as (that which causes RED in JackT). Similarly, on the dispositionalist view it turns out that the original Jack and Jill are not really color inverters (and in fact it is impossible that there be color inverters). Jill is representing a Jill indexed property (red-for-Jill) and Jack is representing a Jack indexed property (red-for-Jack).<sup>20</sup> Thus, this proposal denies existence of intrinsic phenomenal content as well as the possibility of shared perceptual content and true inversion and should be rejected.<sup>21</sup>

A further option is to deny one of the presuppositions of the argument, a presupposition that has been lingering in the background of the debate. This is the assumption that the phenomenal character of an experience determines, without aid of context, a property that is represented. Such phenomenal content is sometimes called *Russellian* (Chalmers 2006; Thompson 2009). On a Russellian theory of phenomenal content, JackT and Jill must be representing blood as being the same color because their experience has the same phenomenal character, and in general any two experiences with the same phenomenal character necessarily represent the same property. It is a natural but unnecessary assumption that phenomenal content is Russellian. Recently Brad Thompson and David Chalmers have pointed out that just as there are Russellian and Fregean accounts of belief content, so there can be more than one account of intrinsic phenomenal content. According to Thompson,

A content is Fregean if it consists of modes of presentation of objects and properties rather than the objects and properties themselves. According to Fregean theories of phenomenal content, the phenomenal content that is shared by any two phenomenally identical experiences is a matter of *how* the world is represented, and need not involve sameness in *what* is represented. . . . For the Fregean, sameness of phenomenal character

need only entail sameness in the mode of presentation. This is compatible with the idea that perceptual experiences do represent particular properties, and that they even have Russellian contents. But this property-involving content will not be, on the Fregean view, *phenomenal* content, since it will be a kind of content that need not be shared by phenomenal duplicates. (Thompson 2009)

So, for example, a Fregean view could maintain that associated with a particular phenomenal character *g*, there is for every subject the following mode of presentation: the property that typically causes *g* in me. The result is that JackT and Jill do have a content in common: the content <Blood is the color property that typically causes RED in me> where RED is the phenomenal experience they both get when looking at things that are red on earth and green on Twin Earth. They will also have a difference in content, however—the (non-dispositional) property they ascribe to blood will be different, since different properties in the world cause RED in each of them. The content they have in common is the phenomenal content. Jack and Jill, however, will have different phenomenal contents, since they have different phenomenal characters, but both are representing veridically: Jack is representing *r*, the color which typically causes GREEN in him, and Jill is representing *r*, the color which typically causes RED in her.

There are several difficulties with the Fregean view.<sup>22</sup> Firstly, this view shares a difficulty with the dispositionalist view. Granting that JackT and Jill can have the same phenomenal characters, they should have the same phenomenal contents. On the Fregean view, they appear to have the same contents—they both have the Fregean content <The color that causes RED in me>. But that is not really a fully determinate condition on reference. It is a description that places a different constraint on reference for different users of the description. Fully articulating the constraint requires that the first person indexical be replaced with its referent. So, for Jill the condition on reference is really <the color that typically causes RED in Jill> and for Jack the condition is really <the color that typically causes RED in Jack>. But these are not really the same contents. So it turns out that the Fregean view has the same problem with inverted spectra as the view that construed colors as dispositions. Jack T and Jill are really not representing things in the same way, and the spectral inverts really aren't true inverts.

It could be objected that though <The color that causes RED in me> is not a fully determinate condition on reference, it is still a condition on reference. There is something in common between utterances of the first person indexical by different subjects that is part of the meaning and content. To object to the Fregean view because the content is not fully determinate might appear to beg the question against the view.

Perhaps it can be granted that there is a level of content shared by different users of the first person indexical, but this only changes the location of the worry. Is it really plausible that phenomenal content is what is shared? The phenomenal

character we are talking about is fully determinate, and when we are imagining two individuals sharing a phenomenal character and thus content we are not imagining them sharing something that is merely determinable. There is, of course, a sense in which someone with a scarlet experience shares content with someone having a vermillion experience—they can perhaps be said to share content on a determinable level. But this is not what we are imagining in the case of JackT and Jill. We are imagining them having the same determinate phenomenal character. We can be pluralists about content and say there is a determinable level of content that the Fregean specifies, but precisely because it is determinable it doesn't seem the right sort of content to capture the intuitions behind the inversion scenarios.<sup>23</sup>

Another problem with the Fregean view is that if it turns out that there is no typical cause of an experience—either because the subject is young, the world is weird, or both—then the perception is not veridical or non-veridical. How does it serve to present the world in such a case? Imagine I am seeing a color for the first time, having an experience caused by that color for the first time, and that the world is such that this will not (or perhaps even could not) be repeated. Does the world really not look to be a certain way? Or does it look to be indeterminate in some manner? This is difficult to reconcile with the first-personal intuitions behind phenomenal content, which is that simply in virtue of the phenomenal character of an experience the subject of the experience represents the world to be a certain way.<sup>24</sup> Since the Fregean view allows that there can be defective modes of presentation that don't actually succeed in representing the world as having a particular property, but that nevertheless have a certain phenomenal feel, the Fregean view seems to contradict this intuition.<sup>25</sup>

Perhaps, though, this is not the way the Fregean view should be interpreted. As I have interpreted it, phenomenological redness is associated with content, the content <the property that causes RED in me>, but perhaps that's unfair. Instead, the Fregean should be taken to say that phenomenological redness is itself the phenomenal content, which is a mode of presentation of whatever the property is that causes RED in the agent. On this reading, JackT and Jill are sharing a fully determinate phenomenal content—as RED is a fully determinate phenomenal state—and when a subject's RED sensation does not have a typical cause, the subject still has a phenomenal content, namely RED. It's just that in such a case RED doesn't present anything.<sup>26</sup>

If this is the view, my previous criticisms misfire, but the cost is that the resulting view is virtually indiscernible from the view according to which there is only a contingent connection between phenomenal states and perceptual content and it thereby falls afoul of the considerations I mounted against that view. The Fregean view only distinguishes itself if there is a level of content associated with the phenomenal state. RED as a phenomenal state by itself is no more of a content than my hand—it is only a content to the extent that it is a semantic feature of the perception by virtue of having truth conditions. The plausible candidate for those truth conditions, as offered by the Fregean we are considering, is just the conditions under which that is perceived has <the

property that typically causes RED in  $S >$  where  $S$  is the subject. So we can consider this a dilemma—either the Fregean view collapses to the ‘mental paint’ view which doesn’t involve intrinsic phenomenal content, or it faces the problems I have suggested.<sup>27</sup>

The biggest worry about the Fregean view (which is made worse, in fact, the more it is similar to views which link content and phenomenal states contingently) is that it seems to flout naive transparency and the cognitive significance that transparency indicates for unreflective perceivers. The idea behind naive transparency was that in the normal course of unreflective perception sensations do not make a cognitive appearance as sensations. That is, their cognitive role does not involve a recognition on behalf of the subject that they are, in fact, sensations and therefore modifications of that subject. On the Fregean view, part of the way the world appears as being involves appeal to subjective states of the subject. That simply isn’t borne out phenomenologically.

The Fregean, of course, might not be trying to capture this phenomenological intuition. Perhaps the Fregean is simply trying to provide an adequate, ‘third-personal’ specification of content according to which there is no requirement that the subject with such contents have even an implicit grasp of the contents. If that is the case, however, the Fregean is simply involved in another project. In fact, however, the Fregean does not seem willing simply to flout the first-person perspective. The modes of presentation invoked in the view are, after all, supposed to capture how things seem to the subject—what is in common between twins that represent different properties. It seems unstable to let this much of the first-person perspective in the door and to exclude other considerations such as transparency and cognitive significance.

To return to the overall dialectic, if the Fregean view is out of the picture, and we accept that the inversion scenarios are possible, then we must conclude that neither Jack nor Jill is representing colors veridically, at least in virtue of their phenomenal content, and thus we aren’t either.<sup>28</sup> The view that naturally results is that at least from the first-person perspective, our perceptual experiences represent properties as in the world that are not in the world. They project the properties responsible for phenomenal character onto the surfaces of objects.

## 6.

There is a rather strong presumption against ‘error theories’ in philosophy. Most of the reasons behind this general presumption, such as those stemming from the principle of charity, work best when applied to a theory which maintains that all of our *beliefs* about a given domain are false.<sup>29</sup> Later in this section I will argue that the error-theorist about perceptual content need not, and should not, claim that this error infects perceptual *beliefs*. If this is correct, then much of the prejudice against error theories does not affect the projectivist view I have favored.

There are, however, important objections that have weight against the projectivist theory in particular. The following three are perhaps the most pressing:

- A) The error-theory maintains that there is a property misascribed by color vision, but there is not a good candidate for what the property is: either it is one instantiated by minds, which is absurd, or it is necessarily uninstantiated, which is equally absurd. (Shoemaker 1990, 1994)
- B) It seems incorrect, uncharitable, and contrary to our ascriptions to say we are always misperceiving colors, etc.
- C) An error theory fails to distinguish between egregious misperceptions, as when someone as a result of bad wiring or as a result of bad lighting concludes that blood looks green, and perceptions we generally take to be veridical.

The first objection is really a demand to flesh out this view of perception. It's one thing to say that we are in error, but until we know what property is really being ascribed by perceptual experiences we haven't really given a story about the content of perception. The second two objections are similar, in that they both fault an error theory for failing to accommodate intuitive judgments about perceptions: in one case that they are usually veridical, and in the other that there are distinctively non-veridical perceptions. I will deal with the first, therefore, by clarifying what sort of property is being misascribed in perception, according to the error theory. I will deal with both of the second by showing how we can recover the truth of most of our intuitive ascriptions of perceptual veridicality and non-veridicality.

There are two natural accounts of the properties ascribed by perceptual experiences: the property ascribed is actually a mental property, the same property that is instantiated by the mind of the ascriber, or the property ascribed is a primitive, nonmental property.<sup>30</sup> Both views seem to have drawbacks. The objection is well voiced by Shoemaker, who calls the previous view 'literal projectivism' and the latter 'figurative projectivism':

As for literal projectivism, I cannot myself make any sense of the idea that any property I perceive as belonging to the surface of the tomato, when I perceive its color, is in fact a property of the experience itself. As for figurative projectivism, it is a mystery to say the least, how the content of our experience can include reference to properties we have never experienced or had any other epistemic access to—properties we know neither 'by acquaintance' nor 'by description' unless we have some sort of nonsensory acquaintance with a Platonic realm of uninstantiated properties. (Shoemaker 1996: 117)

Shoemaker thus finds figurative projectivism implausible because there is no good story as to how we could refer to, and have as part of our perceptual content, properties we have never encountered (and that perhaps are necessarily



uninstantiated). While there is something intuitive about this objection, as it stands it is not conclusive. We can refer to properties that we have never encountered, and perhaps cannot encounter. Perhaps the point is that we can only do so 'by description'. Whether or not this is the case, and whether or not it is a problem if it is the case, depends on further views about how our perceptual states get the content they do. It doesn't seem impossible that one could, by having a hallucination, represent a property one had never experienced—a 'missing' shade of blue, perhaps. Perhaps there are non-question-begging reasons why this isn't possible, or why we couldn't be doing the same thing in perception, but they are not obvious. Nevertheless, it must be granted that something is odd about the view that we project a strange, non-experiential property. When we learn that all that is out there are, for example, surface reflectance properties, we don't tend to form the belief that the properties we were projecting don't exist. We naturally think that we were just mistaken about where they were instantiated. This is admittedly not conclusive, but it seems to sway slightly toward literal projectivism.<sup>31</sup>

Why does Shoemaker find it so implausible that we project our experiential properties onto the surfaces of objects? He continues the critique more powerfully in a later article:

I am looking at a book with a shiny red cover. The property I experience its surface as having, when I see it to be red, is one that I can only conceive of as belonging to things that are spatially extended. How could that property belong to an experience or sensation? Remember that an experience is an experiencing, an entity that is 'adjectival on' the subject of experience. It seems no more intelligible to suppose that a property of such an entity is experienced as a property of extended material objects than it is to suppose that a property of a number, such as being prime or being even, is experienced as a property of material things. (Shoemaker 1996, 250–1)

The gist of Shoemaker's objection is that not only does the error theory commit us to persistent non-veridicality, but the error involved is so egregious as to be nonsensical.

The absurdity Shoemaker has in mind involves a sort of category mistake. It's one thing to say that a brown thing is blue, but another to say that a number is jolly. But for a category mistake to be made, it seems categories need to be deployed, and at the level of phenomenal content we are not at that level of sophistication. The error theorist need not ascribe to anyone the belief that their mental states, thought of as mental states, are spread across the surfaces of object. The projectivist is committed to saying that there is a quality, which in fact only characterizes mental states, that appears to characterize things in the world. This does not imply that anyone thinks that something falling under a category of 'way of experiencing' characterizes an insentient book. It is, however, the case that it seems as though a property, which we as theorists know is adjectival on the mind, is modifying the world.<sup>32</sup>

Perhaps this response mistakes the import of Shoemaker's objection. Perhaps Shoemaker isn't remarking on the absurdity of the category mistake made by the perceiver in perception—he can grant that the naive perceiver isn't employing categories. What he finds impossible to understand is the theoretical claim—made by the theorist who is using concepts—that the perceiver is attributing to a book a property that is in fact an experience, dependent ontologically on an experiencer.<sup>33</sup>

What is it that makes this theoretical claim unintelligible, however? It doesn't become unintelligible simply because we as theorists who know better cannot see how one could reasonably think of something mental and 'a way of appearing' as being spread across the surface of a book. If the theorist can explain why something in one ontological category can appear to be in another, it seems the error is intelligible enough. Consider a parallel case.

Stories abound about the audiences in the early days of cinema leaping out of their seats because of action happening on the screen. When watching the film *Rough Sea at Dover*, for example, theater patrons were afraid of getting splashed. According to a reporter, 'Wave after wave came tumbling on the sand, and as they struck, broke into tiny floods just like the real thing. Some people in the front row seemed to be afraid they were going to get wet, and looked to see where they could run, in case the waves came too close.'<sup>34</sup> From one perspective this error is absurd. The theatergoers are attributing to patterns of light on a two-dimensional screen the property of leaping out and making them wet. How could the patrons be representing light as wet, and the patterns on the screen as leaping? These are metaphysical absurdities. But we as theorists of the illusion do not need to make sense of the metaphysical impossibilities involved here. We must only explain how things could seem the way they do to the people under the grip of the illusion. Of course the way to do that is not to say that their perceptual system is representing <this light is going to get me wet> or anything like that. Such a mistake would be hard to understand. But it might be right, though perhaps misleading, to say that their systems are representing the light as water that is about to get them wet, when their content is something more like <that stuff is going to get me wet! > As long as we can accept that what is actually light can look like water, there is no real absurdity from the perspective of the theorist of illusion.

It seems we are in a similar boat with the illusion that red experiences are covering books. As theorists, when we say people are making this error we are not saying that they have the content <the red experience is on my book>. We are saying something that they have a content more like <my book has that feature>. There is no absurdity so long as we can accept that things that are as a matter of fact mind-dependent experiences can fail to appear to be such. As a matter of fact, we should be quite used to the idea that phenomenal experiences can be presented in ways that are really quite contrary to their natures. We make attributions as theorists, both to other theorists and to unreflective subjects, that posit similar errors. We say that dualists make the mistake of thinking that something that is as a matter of fact a brain state can exist without brains—or

can fail to exist despite the fact that the brain state exists. Idealists claim that we believe of experiential states that they are material. Shoemaker himself accuses philosophers of thinking of sensations as objects when they are themselves adjectival on the mind (Shoemaker 1996: 209). It is not far fetched to think, in fact, that both the error of naive perceiver and the errors of the dualist stem from the same source: the fact that our epistemic relationship to our sensations is very peculiar and does not reveal their nature (Conee 1994; Hill 1991; Howell 2013; Papineau 2002; Pereboom 2011) This fact, which is widely accepted by philosophers working on consciousness, ought to be enough to satisfy us that the error we are attributing to naive perceivers is not absurd.

Even if we are convinced that the error in error theories is not flat-out incomprehensible, there remains the objections voiced in B and C. It seems uncharitable to say that we are consistently misrepresenting the world. This falsifies almost all of our ascriptions of perceptual veridicality. In addition, it makes a hash of our distinction between normal perceptions—which are the result of proper functioning perceptual systems—and real misperceptions such as when I suddenly see grass as the color I have previously seen in ripe tomatoes.

In response, a distinction should be made between a sensation's being non-veridical and its being misleading. Just because all color sensations are non-veridical does not mean they are equally misleading.<sup>35</sup> Consider the appearance of a stick partially submerged in water. It seems reasonable to say that the stick does look a certain way—it looks bent. This appearance is in error, and systematically so. But just because there is an error here doesn't mean that one couldn't have *worse* perceptual errors under the circumstances. If, due to some mental mishap, the stick submerged in water looked as though it turned in on itself like a spiral, that would be inapt, and a much worse representation as far as the subject is concerned. Not all misrepresentations are on par. In some sense, a stick is supposed to look bent when partially submerged in water—even if it is not really bent—and it is not supposed to look like it spirals in on itself. Ironically, if the stick appears straight, the appearance is veridical but not apt, because it is in some sense supposed to look bent.

What is the sense of 'supposed to look like F' here? One appealing suggestion is that something *x* being *G* is supposed to look like *F* for a subject *S* in circumstances *C* iff *x* typically looks like *F* to *S* when it is *G* in *C*. The appearance that *x* is *F* can be apt, even if *x* is not in fact *F*. It is apt because it offers a guide by which the subject can keep track of the properties in the world, even if the way the properties appear is not how they are. The fuel gauge in my car might systematically lie, telling me that there is an eighth of a tank more gas than in fact there is. If it has always been that way, however, and I am used to judging in accordance with that appearance, adjusting for the error, this representation can help me keep track of the way things are in the world. Ironically, perhaps, I am more likely to be fooled into forming a false belief if the gauge suddenly becomes accurate—I will almost certainly begin forming false judgments that the car has an eighth of a tank less gas than it does.

An impression can be apt due to its reliable correlation with a property in the external world even if the representational content of the impression is not veridical. The distinction between apt and inapt sensations helps us address objection C. The error theorist can accommodate the distinction between normal perceptions (which the error theory deems non-veridical) and egregious misperceptions (which result from internal distortions in the agent, unusual situations, etc.). Both are non-veridical, but the normal perceptions are apt and the misperceptions are inapt.

This still leaves objection B, which is that it is uncharitable to say that all of our experiences are non-veridical. In fact, there are really four distinct ways in which an error theory might be considered uncharitable. One (we can call this objection B1) is that it is uncharitable to the perceiver to say he always perceives non-veridically. The second (B2) is that it is uncharitable to ascribe consistently false perceptual beliefs to the perceiver. But it is not just to the perceiver that the error theory might be thought to be uncharitable. The error theory also renders false many judgments made about perceivers. So this generates two further worries about charity: (B3) is that it is uncharitable to those who judge the perceiver's experiences to be veridical—a great deal of their judgments about veridicality turn out to be false. And there is also (B4) it is uncharitable to those who judge the perceiver's perceptual beliefs to be true.

I am not always sure what to make of claims about charity. If there is reason to believe that there is a persuasive illusion in a subject's environment, and the subject shows no awareness of the illusion, is it really more charitable to say that the subject wasn't fooled? I, for one, have trouble taking this sort of consideration too seriously. There is, however, something to the idea that it is unfair to convict someone of a mistake in certain cases. So, for example, when a subject says 'x is green' it might be fair to say that it only looks that way because of the lighting, but it doesn't seem a fair criticism that no, it isn't green (in the way they think) because nothing is. A wholesale condemnation of almost all common perceptual beliefs and beliefs about perceptions seems to ignore the fact that even if something is wrong about the beliefs, much more seems to be right about them.

One tempting answer to this concern is to distinguish between the truth conditions of beliefs and attributions and the appropriateness conditions. On this view, the perceptions, perceptual beliefs, and the attributions of veridicality are all strictly speaking false or non-veridical. But what determines the appropriateness of these ascriptions is not the veridicality of the perceptions in question but their aptness. To object to someone's claim that 'Roger is right about the color of the book' is to make it seem as though one thinks Roger's perceptions are inapt or particularly off, when in fact things are appearing how they are 'supposed to appear'. This is a similar move to that made by the epistemological skeptic who says that though no one knows anything, saying of the history professor that she doesn't know who won the battle of Waterloo implies something untrue—that the professor has made some particularly egregious mistake, rather than sharing with all of us the inability to rule out evil demon

scenarios. The fact that denials of veridicality (or knowledge) have false implications might explain why they seem uncharitable, even if they are true.

Though this is a start of a response to the objections about charity, we can do better. The error theory we have presented is a first-personal theory of perceptions. It is not a theory of perceptual beliefs or of judgments about perceptions and perceptual beliefs. The best move for an error theory, I suggest, is to leave open the possibility that though perceptions are systematically non-veridical, this non-veridicality does not lead to the falsity of perceptual beliefs, or the falsity of judgments about perceptions and perceptual beliefs.

Consider the beliefs in B3 and B4—beliefs or judgments about the veridicality of the perceptions of others and the truth of other people's perceptual beliefs. The sourest taste generated by the error theory is plausibly caused by the fact that it seems to render all these beliefs false. But it is open to our error theory, driven as it is by the first person perspective, to claim that our judgments about the perceptions of others are not driven by the same considerations. Quite plausibly, what we care about in third personal judgments about perceptions is aptness. We are concerned about whether or not the subjects can be relied upon to track properties in their environment and whether or not they are tracking the same properties everyone else is tracking. These judgments play the role of coordinating action and behavior—whether in the form of hiring interior decorators or operating on glaucoma—without any need to be particularly sensitive to the way things appear from the first person perspective. There is no reason to insist that the truth conditions of veridicality attributions, then, depend upon the veridicality of phenomenal content. Phenomenal content and its veridicality just isn't particularly important at this level. Saying that someone's experience is veridical is quite plausibly to be taken to comment on something else entirely—namely the aptness of the content.

If this is plausible, we have an answer to B3 and B4. The error theory does not say that such beliefs are false, and there is a plausible story according to which they are true. What about B1 and B2? B1, I think, can largely be set aside if we can handle B2. That is, it doesn't seem uncharitable to say that a subject's perceptual appearances are not veridical if this doesn't give rise to systematically false beliefs. The fact that things appear a certain way to the subject is something for which he is not responsible, and if we acknowledge that at the more sophisticated cognitive level there are no mistakes, it doesn't seem we are being uncharitable. (Compare—is it uncharitable to say that we always misperceive oars when they are partially submerged?)

In fact, there are compelling reasons to claim that a subject's perceptual beliefs are veridical even if phenomenal content systematically misrepresents. One simple reason is that his beliefs about his own perceptions, and his beliefs formed on the basis of those perceptions, play the same sort of role as the third personal beliefs mentioned in B3 and B4. Consider a subject's reports about his own perceptually based belief—'x is green'. What purpose does this serve—the purpose of reporting on the way things appear or the purpose of reporting about the properties of things in the world? Clearly the latter. What tribunal is faced

(and is meant to be faced) by such a statement? It is the confirmation by others (and by future experience) that the subject is tracking the correct properties in his environment. The subject's belief 'x is green' and his observer's belief 'S's belief that x is green is true' are intimately connected; more intimately connected than the subject's belief and his perceptual contents. If there is a reason for thinking that the observer belief is governed by the subject's success as tracking properties in his environment and not by the veridicality of phenomenal appearances, there is reason for thinking the subject's belief is governed in the same way.

On this picture—which is not mandatory for the error theorist, but is open to him—the subject's perceptual beliefs do not inherit the error from his appearances, but instead inherit the roles, purposes and adequacy of the concepts and beliefs of the intersubjective social world of which he is a part.<sup>36</sup> This would sufficiently rebuff the charges of lack of charity while preserving the first-personal, phenomenological core of the view.<sup>37</sup>

This view, which gives beliefs different contents from phenomenal states, has an advantage over a similar view which avoids attributing rampant error by claiming perceptual content itself has two levels, one veridical and one non-veridical (Chalmers 2006). On such a view, there could be one part of phenomenal content that is erroneous, but that there is also a Fregean content which is veridical. So, a subject's perception has at one and the same time the phenomenal content that  $\langle x \text{ is } G^* \text{ish} \rangle$  where  $G^*$  is a sensation, and  $\langle x \text{ has the sensation that typically causes } G^* \text{ in me} \rangle$ . Such a view could be said to accept split phenomenal content.

Split contents are problematic. For one thing, the view that splits contents seems to inherit the problems with the Fregean view, most importantly the neglect of naive transparency. This neglect seems particularly problematic for the split view, because one and the same state has a content that recognizes sensations as sensations and another that doesn't. This makes for an odd picture of the subject's worldview—does he recognize his sensations or not? It is only plausible to say that a subject's experience is representing the external world as having a quality that only sensations can have because at the level of naive perceptions the subject is not sensitive to his sensations *qua* sensations. But if a subject is in the first instance perceiving things as being the causes of certain sensations, then it is odd to say that by having the very same experience the subject is projecting his sensations onto the surfaces of objects. This split view therefore attributes to the subject a basic sort of incoherence. This is uncharitable in any case, but it certainly seems to flout the first-personal intuitions which originally motivated the idea of phenomenal content.

## 7.

In the unreflective, first-personal point of view, experience is transparent or diaphanous in the sense that we are not aware of our sensations or experiences as such but are riveted to the properties of the external world. I have argued,

however, that this does not entail that we are not directly aware of our sensations, even in this unreflective mode. If we are to respect the point of view of the subject, however, this transparency should be taken into consideration when developing a theory of the content of perceptual experience. From the first-person point of view, there is, I have argued, a very tight connection between the phenomenal character of experiences and the phenomenal content of experience—so tight, that the phenomenal character seems intrinsically presentational in nature. These observations, along with considerations of color-inversion scenarios, suggest that the best view of phenomenal content is projectivist.

The resulting picture is that we are, even unreflectively, directly aware of our experiences, but without being aware of them as our experiences. Instead, they appear to paint and pervade the world around us. This is a significant representational error, but it need not result in systematic error in perceptual belief. Perceptual beliefs, and beliefs about the veridicality of perceptions, do not simply appropriate the veridicality conditions of perceptual appearances making them truth conditions for the belief. Rather, perceptual beliefs are true if they attribute the property to objects that typically causes those perceptual appearances in the subject.

Of course this view will not appeal to everyone. Even if the error is limited to perceptual content, the distaste of embracing an error theory will lead many to look for other options. There are always many places to get off the boat, but I would suggest that for these philosophers, the arguments of this paper can be recast as follows: if there is intrinsic phenomenal content, our perceptions would never be veridical. This is unacceptable. So, there must not be intrinsic phenomenal content. In other words, it seems to me that the place to cut the line is at the very beginning. Once one accepts that there is phenomenal content, content that supervenes upon what it's like for the subject, it seems that things appear differently to Jack and Jill and the same to Jack and JackT. The only plausible ways to avoid an error theory at this point is to loosen our commitment to the first-person perspective and to deny that a subject has particularly good access to the representational features of his perceptions. To do this, however, seems to undermine our motivations for saying there is phenomenal content in the first place. Unless principled reasons can be provided for abandoning the first-personal stance at a particular point, this position seems to be unstable. We seem to be in the position of either being guided by how things appear to the subject, or theorizing from the outside. To do the former risks landing us in systematic error, and to do the latter makes issues of transparency, or how things phenomenally seem to the subject, largely irrelevant.<sup>38</sup>

*Robert J. Howell*  
*Department of Philosophy*  
*Southern Methodist University*  
*USA*  
*rhowell@smu.edu*

## NOTES

<sup>1</sup> Many of Sartre's comments on the emotions also sound surprisingly like Harman's famous claims about transparency. See, for example, Sartre (1993b: 51).

<sup>2</sup> This transparency thesis should not be confused with a weaker thesis, that one is not *usually* directly aware of one's own experiences, or that it is *difficult* to be aware of them. For insight into this area, I am indebted to Kind (2003).

<sup>3</sup> See for example Shoemaker (1996).

<sup>4</sup> There might, perhaps, be a sense of 'direct awareness' such that only sense data theorists believe we are directly aware of our sensations. That is surely a technical sense as well.

<sup>5</sup> There might be notions of perception according to which I would count as perceiving the rosiness of the lenses, but that is not a problem for what I say here. I am merely claiming that if one wants to reserve 'perception' for the state that meets a specific set of conditions—such as those in Shoemaker (1996) pp. 204–206, there is another more liberal notion of awareness which might be more relevant and more intuitive.

<sup>6</sup> See, for example, the debate between Peacocke (1983), Block (1996) and Tye's responses in Tye (2000).

<sup>7</sup> Tye discusses the afterimage case, and other similar cases, in Tye (2000) pp. 83–88.

<sup>8</sup> Fernandez (2003) and Byrne (2005) have flirted with this idea, but see Zimmerman (2004) and Gertler (forthcoming) for replies.

<sup>9</sup> If one is worried that one might necessarily be able to be aware of a television set without being aware of what it represents, one can perhaps imagine virtual reality goggles or something of the sort.

<sup>10</sup> Despite what this example suggests, the experience first model should not necessarily be thought to be committed to sense data.

<sup>11</sup> One is not directly aware of the pixels simply because they are *causal intermediaries*. There is an intuitive difference between the role the pixels play in one's representation and the role the intervening air plays.

<sup>12</sup> My terminology borrows from, but is not the same as, that of Siewart (1998), Chalmers (2006), Thompson (2009), and Shoemaker (1996).

<sup>13</sup> Chalmers and Thompson, among others, define phenomenal content as content that supervenes on phenomenal character. This seems to exclude the possibility of the contingent connection view. Thanks to Jesse Prinz for urging the change in terminology.

<sup>14</sup> Note that such a view need not say that intrinsic phenomenal content is all there is: a hybrid view is possible. This will be close to, but not identical to, the view Chalmers (2006) adopts, and similar to but different yet again from the view I will adopt below. Such a view also need not be representationalist, in the sense dominant in the literature, according to which phenomenal character supervenes upon phenomenal content.

<sup>15</sup> The arguments in the previous paragraph owe a good deal to Siewert (1998).

<sup>16</sup> I am not arguing from an illicit conflation of 'seeming' and 'seeming to be', as Chalmers (2006) suggests some might be guilty of. The connection is instead a conclusion of the considerations in the previous paragraph.

<sup>17</sup> There are two options available. One is to say that GREEN represents 'F or G' and this is in fact shared by a and b. But this is a weak fix. The sense in which it seems to me the two things have something in common is precisely non-disjunctive. The other is to say the two share a dispositional property, to cause GREEN in me. But it seems to be part of the transparency intuition that at least in the first instance, GREEN doesn't make



an appearance separate from the items in the world that I am perceiving. So this view would have to flout that intuition.

<sup>18</sup> Though see Tye (2000) and (2002) for a contrary opinion.

<sup>19</sup> For simplicity's sake I'm supposing a simple dispositionalism, but more sophisticated relationalist views could be adopted. They will, however, have the same difficulties in this context.

<sup>20</sup> It's worth noting that one way out of the Jack/Jill example, indexing the colors to perceiver types instead of individual perceivers, fails in the case of JackT and JillT since they are different perceiver types.

<sup>21</sup> These arguments owe deeply to Chalmers (2006) and Thompson (2007).

<sup>22</sup> I say 'the' Fregean view, though there really could be many, depending on what the mode of presentation associated with phenomenal content is. I focus on this case not only because it is the one defended in the literature, but because I suspect other Fregean views will have similar or corresponding difficulties.

<sup>23</sup> Thanks to Philippe Chuard and Brad Thompson for debate on this objection.

<sup>24</sup> A referee suggests that in such a case the experience would still have phenomenal content. This is true in a sense, since the experience is still associated with conditions that determine reference—they just don't determine a particular reference in this case. The argument here, though, is that this generates the counterintuitive result that this content really doesn't present the world to be a certain way.

<sup>25</sup> In fact, it might even be *likely* that there is no typical cause of a particular phenomenal experience. Given that the experience created in me by seeing a red 27 in sunlight might be the same as the experience of seeing the much brighter red 68 in shade, there is probably not a simple causal match between phenomenal characters and color properties that the Fregean view seems to suppose. (Though by making the view holistic, this might be able to be handled. See Thompson (2006). Introducing holism at this level, though, might introduce its own implausibilities.)

<sup>26</sup> Thanks to an anonymous referee at this journal for pushing this objection. Thompson has indicated in conversation that this is not the route he prefers, but it is nevertheless a possible position.

<sup>27</sup> The problem remains even if one resists the somewhat descriptivist gloss I am giving to the content. Since the phenomenal content is the phenomenal analogue to the 'primary intensions' of Chalmers (1996), they are to be associated with a function that takes *centered* possible worlds as inputs, where the centers are determined (among other things) by the individuals who have those contents. See, for example, Chalmers (2010) p. 259.

<sup>28</sup> Of course some philosophers deny the conceivability and/or the possibility of the inversion scenarios. See, for example, Tye (2000) and (2002).

<sup>29</sup> Even then, there is real doubt that the presumption is justified. For a nice discussion of this issue, see Daly and Liggins (2010).

<sup>30</sup> Boghossian and Velleman (1989) suggest the former, Thau (2002) and McGinn (1996) the latter. (Thau's view is actually more complex, involving a double attribution, which in the end makes his view similar to the one proposed here.)

<sup>31</sup> Whether or not the property being projected is mental or is some primitive property is not crucial for what I am arguing here. I don't yet see conclusive arguments that point either way, but it does seem that if there is a primitive property, it had better have a rather close relation to the properties responsible for the phenomenal characters of experiences.

<sup>32</sup> Of course there is puzzlement as to how something like this can modify the world, but there is equally well a puzzle as to how anything such as qualia could be nothing more than brain states either. To some eyes, it will seem just as absurd to make such a

claim, so it's not clear that this is an area where our sense of conceptual outrage should carry much weight.

<sup>33</sup> Thanks to a referee from this journal for forcing me to deal with this version of Shoemaker's complaint.

<sup>34</sup> Quoted from <http://www.precinemahistory.net/1895.htm>

<sup>35</sup> Cf. Sosa (2007) and much of his work in the preceding decade.

<sup>36</sup> What I am saying here obviously draws a lot from Burge (1979). One way of putting my argument is that social externalist considerations prevail when it comes to beliefs, but not when it comes to appearances.

<sup>37</sup> What then is the content of the various beliefs? It seems plausible that while S's perception (of green grass, say) has the content of <that is G\*ish> where G\* is the visual sensation grass typically causes in the subject, his belief is <that is G> where G is the property that typically causes G\* in the subject. When his observer says 'Roger is correctly perceiving that the grass is green' her belief has the content <Roger is correctly perceiving that the grass is G>, not <Roger is correctly perceiving that the grass is G\*ish>.

<sup>38</sup> Thanks to groups at NYU and SMU for helpful comments on an earlier draft of this paper. In particular, thanks to Katilin Balog, Jacob Berger, Ned Block, Richard Brown, David Chalmers, Philippe Chuard, Imogen Dickie, Justin Fisher, Farid Masrour, Jesse Prinz, Jim Prior, Jon Simon, and Brad Thompson.

## REFERENCES

- Block, N. (1990), 'Inverted Earth', *Philosophical Perspectives*, 4: 53–79.
- (1996), 'Mental Paint', in M. Hahn and B. Ramberg (eds) *Reflections and Replies: Essays on the Philosophy of Tyler Burge*. Cambridge, MA: MIT Press, 165–200.
- (2005), 'Bodily Sensations as an Obstacle for Representationism', in M. Aydede (ed.) *Pain: New Essays on Its Nature and the Methodology of Its Study*. Cambridge, MA: Bradford Book/MIT Press, 137–42.
- Boghossian, P. A. and Velleman, J. D. (1989), 'Color as a Secondary Quality', *Mind*, 98: 81–103.
- Brewer, B. (2006), 'Perception and Content', *European Journal of Philosophy*, 14: 165–81.
- Brown, D. H. (2009), 'Indirect Perceptual Realism and Demonstratives', *Philosophical Studies*, 145: 377–94.
- Burge, T. (1979), 'Individualism and the Mental', *Midwest Studies in Philosophy*, 4: 73–122.
- Byrne, A. (2005), 'Introspection', *Philosophical Topics*, 33: 79–104.
- Chalmers, D. (1996), *The Conscious Mind*. New York: Oxford University Press.
- Chalmers, D. J. (2006), 'Perception and the Fall From Eden', in T. S. Gendler and J. Hawthorne (eds) *Perceptual Experience*. Oxford: Oxford University Press, 49–125.
- Conee, E. (1994), 'Phenomenal Knowledge', *Australasian Journal of Philosophy*, 72: 136–50.
- Daly, C. and Liggins, D. (2010), 'In Defense of Error Theory', *Philosophical Studies*, 149: 202–30.
- Dretske, F. (1995), *Naturalizing the Mind*. Cambridge, MA: MIT Press.
- Fernandez, J. (2003), 'Privilege Access Naturalized', *Philosophical Quarterly*, 53: 352–72.
- Gertler, B. (forthcoming), 'Self Knowledge and the Transparency of Belief', in A. Hatziioyis (ed.) *Self-Knowledge*. Oxford: Oxford University Press, 125–45.
- Harman, G. (1990), 'The Intrinsic Quality of Experience', *Philosophical Perspectives*, 4: 31–52.

- Hill, C. S. (1991), *Sensations: A Defense of Type Materialism*. Cambridge: Cambridge University Press.
- Howell, R. (2013), *Consciousness and the Limits of Objectivity*. Oxford: Oxford University Press.
- Jackson, F. (1977), *Perception: A Representative Theory*. Cambridge: Cambridge University Press.
- Kind, A. (2003), 'So What's So Transparent about Transparency?', *Philosophical Studies*, 115: 225–44.
- McGinn, C. (1996), 'Another Look at Color', *The Journal of Philosophy*, 93: 537–53.
- Moore, G. E. (1903), 'The Refutation of Idealism', *Mind*, 12: 433–53.
- Papineau, D. (2002), *Thinking About Consciousness*. Oxford: Oxford University Press.
- Peacocke, C. (1983), *Sense and Content: Experience, Thought, and Their Relations*. Oxford: Oxford University Press.
- Pereboom, D. (2011), *Consciousness and the Prospects of Physicalism*. Oxford: Oxford University Press.
- Sartre, J.-P. (1957), *The Transcendence of the Ego: An Existentialist Theory of Consciousness*. New York: Octagon Books.
- (1993b), *The Emotions*. Trans. Bernard Frechtman. New York: Carol Publishing Group.
- Shoemaker, S. (1990), 'Qualities and Qualia: What's in the Mind?', *Philosophy and Phenomenological Research Supplement*, 50(Suppl): 109–31.
- (1994), 'Lecture III: The Phenomenal Character of Experience—Self Knowledge and Inner Sense', *Philosophy and Phenomenological Research*, 54: 291–314.
- (1996), *The First-Person Perspective and Other Essays*. Cambridge: Cambridge University Press.
- Siewert, C. (1998), *The Significance of Consciousness*. Oxford: Oxford University Press.
- Sosa, E. (2007), *A Virtue Epistemology: Apt Belief and Reflective Knowledge*. Oxford: Oxford University Press.
- Stoljar, D. (2004), 'The Argument From Diaphanousness', in M. Escurdia, R. J. Stainton and C. D. Viger (eds) *Language, Mind and World: Special Issue of the Canadian Journal of Philosophy*, vol. 30. New York: University of Alberta Press, 341–90. Supplement.
- Thau, M. (2002), *Consciousness and Cognition*. New York: Oxford University Press.
- Thompson, B. (2006), 'Color Constancy and Russellian Representationalism', *Australasian Journal of Philosophy*, 84: 75–94.
- (2007), 'Shoemaker on Phenomenal Content', *Philosophical Studies*, 135: 307–34.
- (2009), 'Senses for Senses', *Australasian Journal of Philosophy*, 87: 99–2013.
- Tye, M. (2000), *Consciousness, Color, and Content*. vol. 113. Cambridge: MIT Press.
- (2002), 'Visual qualia and visual content revisited', in D. Chalmers (ed.) *Philosophy of Mind: Classical and Contemporary Readings*. Oxford: Oxford University Press, 447–56.
- Zimmerman, A. (2004), 'Unnatural Access', *Philosophical Quarterly*, 54: 435–8.