

# Transparency, intentionalism, and the nature of perceptual content

May 15, 2007

Few acts of introspection have been as philosophically influential as G. E. Moore's examination of his experience of a blue surface:

“...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.”<sup>1</sup>

Moore's idea here is often summarized by saying that experience is transparent: when we try to examine the features of an experience, we end up ‘looking through’ the experience and examining features of what the experience is an experience of.<sup>2</sup> Though there is widespread agreement that the transparency of experience shows something important about perception, there is little agreement about what it shows.

Many have argued that we can use it to decide questions about the objects of experience — whether they are sense data, propositions of some sort, or external particulars and their properties. I agree that there is something important to be learned from Moore's observation: but I think that it concerns the contents, not the objects, of perception.

These two topics are not the same. To ask whether perceptual experiences have objects is to ask whether the having of a perceptual experience is a matter of instantiating a certain monadic property or of bearing a certain relation — call it ‘perceptual awareness’ — to something. To ask what the *objects* of perception are is to ask, given that the latter view of the metaphysics of perception is correct, what sorts of things the second relation of the perceptual awareness relation can be. To ask about the nature of the *contents* of perceptions is to ask, given that there are ways that perception represents the world as being — what are these ‘ways’?

On certain views, there may well be connections between the objects and contents of perceptual experiences. But it is not obvious that everyone must recognize a close link here. An adverbialist thinks of the having of a perceptual experience as instantiation of a certain monadic property, and therefore denies that there are such things as the objects of perception, in the above sense. Must the adverbialist therefore deny that perceptions have contents — that perceptions sometimes represent a table as brown, or

---

<sup>1</sup>Moore (1903), 450.

<sup>2</sup>The use of ‘transparency’ in this context is also due to Moore: “that which makes the sensation of blue a mental fact seems to escape us; it seems, if I may use a metaphor, to be transparent — we look through it and see nothing but the blue . . .” (Moore (1903), 446).

a sound as coming from the left? Surely not. And it also seems coherent to believe that perceptions have both contents and objects while declining to identify the two. This point is important, because the views that perceptual experiences have contents and that they are fundamentally relations to propositions are often conflated, and skepticism about the latter position often leads to skepticism about the former.

There are, to be sure, other sources of skepticism about the thesis that perceptions have contents.<sup>3</sup> Often, I think, arguments for skepticism about perceptual content are based on dispensable additions to that thesis and, for this reason, I think that many of the claims which follow about the contents of perception could be stated in ways which opponents of the idea that perceptions represent the world would find acceptable. But arguing that point is beyond the scope of this paper and so, for that reason, in what follows I will simply assume that it makes sense to talk about the contents of perceptual experience.

My strategy in what follows will be to give an interpretation of the phenomenon to which Moore is calling our attention in the above quote, and to argue that this phenomenon provides the basis of arguments both for intentionalism — understood as the view that there is a necessary connection between perceptual content and perceptual phenomenology — and for the view that the contents of perceptual experiences are Russellian propositions. While each of these views is popular, there are apparent tensions between them, and some have thought that their combination is unstable.<sup>4</sup> In the second half of the paper, I'll respond to these worries by arguing that Russellianism is consistent with intentionalism, that their conjunction is consistent both with internalism about phenomenal character and externalism about perceptual content, and that the resulting view receives independent support from the relationship between hallucination and thought.

### 1 *From transparency to intentionalism*

Given the assumption that experiences have contents, what does Moore's point show about them? One thing that it shows is that there is a close link between the *content* of a perceptual experience and its *phenomenology*, or *phenomenal character*. Both of these latter terms have been given a variety of interpretations. But for now let's simply identify the phenomenology or phenomenal character of an experience with 'what it's like' to have the experience. This is hardly a perfectly clear definition; but it's hard to deny that we have some grip on the idea of two experiences seeming the same, or being indistinguishable to, a perceiver. Given this way of understanding the terms, it is clear that the contents of experiences are at least notionally distinct from their phenomenologies. One might have the view that phenomenology can be explained in terms of content, or vice versa; but these would be substantive views, asserting a link between conceptually distinct aspects of perceptual experiences.

To see the what Moore's point about transparency shows about the relationship between content and phenomenology, consider the following oft-quoted gloss on Moore's idea, from 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

---

<sup>3</sup>For examples, see Chapter 6 of Campbell (2002), Travis (2004), Alston (2005), Brewer (2006), and the discussion of semantic and stative conceptions of representation in Martin (2002).

<sup>4</sup>See, for example, Chalmers (2004) and Caplan and Schroeder (forthcoming).

features of anything as intrinsic features of her experiences. And that is true of you too. There is nothing special about Eloise's visual experience. When you see a tree, you do not experience any features as intrinsic features of your experience. Look at a tree and try to turn your attention to intrinsic features of your visual experience. I predict you will find that the only features there to turn your attention to will be features of the presented tree. ..."<sup>5</sup>

Harman's idea is that introspection reveals nothing which is not an aspect of how the scene before one is presented as being; we notice only the objects that are represented as being in one's environment, and the properties those objects are represented as having. I think (unoriginally) that this claim is phenomenologically plausible, and that for this reason it provides a good starting place for theorizing about perception.

Harman's idea can be stated as follows:

#### Transparency

Nothing is available to introspection other than the objects represented as in one's environment, and the properties they are represented as having.

This is a principle about what we notice when we introspect our experiences. It seems to entail the following principle about introspectable differences between experiences:

#### Transparency/Difference Principle

If there is an introspectable difference between two experiences, then there is a difference in the objects and properties those two experiences represent as in one's environment.

If the only objects of introspection are the objects and properties presented as in the environment of the perceiver, then any introspectable difference between two experiences must correspond to a difference in the represented objects and properties.<sup>6</sup>

This gives us the makings of an argument that there is a necessary connection between content and phenomenology. If we think of the phenomenal character of an experience as 'what it's like' to have the experience, then it is plausible that any difference in phenomenal character between two experiences must be introspectable (even if not actually introspected). Any difference between the two experiences which was in principle not noticeable by the subject would not be a difference in what it is like for the subject to have those experiences. This link between phenomenal character and introspection suggests the following argument:

---

<sup>5</sup>Harman (1990), p. 667.

<sup>6</sup>It is worth noting that there is a sense in which not much hangs on this inference, for we can give the same kind of direct argument from introspection for the Transparency/Difference Principle as Moore and Harman give for Transparency. Suppose that introspection reveals a difference between your experiences of a flat surface over time. Suppose (for simplicity) that the only introspectable aspect of your experience of the surface is its color. Can there be an introspectable difference between two experiences of the wall which is not a difference in the color that the wall seems to have?

There are complications here involving the connection between visual phenomenology and shifts in attention, which I ignore for present purposes. I discuss these in ?.

1. If two experiences differ in phenomenal character, there is an introspectable difference between them.
  2. If there is an introspectable difference between two experiences, then there is a difference in the objects and properties those two experiences represent as in one's environment. (Transparency/Difference Principle)
  3. If there is a difference in the objects and properties two experiences represent as in one's environment, there is a difference in the content of the two experiences.
- 
- C. If two experiences differ in phenomenal character, they differ in content.

The conclusion of this argument deserves the name 'minimal intentionalism.'<sup>7</sup> This version of intentionalism is minimal in two (related) respects, each of which will be important in what follows: it is a claim about the supervenience of phenomenal character on content, rather than the identity of phenomenal character with content; and it states that any change in phenomenal character must be accompanied by a change in content, but not the converse claim that any change in content must show up as a change in phenomenal character.<sup>8</sup>

The third premise of this argument for minimal intentionalism will hold on any plausible view of content; but one might have a reservation about the first premise. Perhaps if two consecutive experiences of a single subject differ in phenomenal character, there is a clear sense in which there must be some introspectable difference between them. But if the two experiences are separated by a period of years, for example, are we so sure that a difference in phenomenal character must be introspectable? And if we are comparing the experiences of two different subjects, in which case there is no subject to do the introspecting, what does it even mean to say that a difference in phenomenal character between two experiences must be introspectable?

These questions bring out the fact that there are three importantly different versions of the first premise of this argument, and so also three corresponding versions of minimal intentionalism, some more minimal than others:

Intrapersonal, time-restricted minimal intentionalism

If two consecutive experiences of a single subject differ in phenomenal character, then they differ in content.

Intrapersonal, time-unrestricted minimal intentionalism

If two experiences of a single subject (whether consecutive or not) differ in phenomenal character, then they differ in content.

Interpersonal minimal intentionalism

If two experiences (whether of a single subject or two subjects) differ in phenomenal character, then they differ in content.

---

<sup>7</sup>For related arguments for intentionalism, see Harman (1990), §3 of Byrne (2001), Tye (2002).

<sup>8</sup>Minimal intentionalism is thus a 'non-reductive' version of intentionalism, in the sense of Siewert (2004). An important issue which I'm ignoring for present purposes is the question of whether the statement of intentionalism should contain a 'of the same sense modality' qualifier. At first glance, it looks as though the argument from Transparency should support an unqualified version of intentionalism if it works at all.

Whether the above argument supports anything stronger than intrapersonal, time-restricted intentionalism depends in part on whether we can make sense of comparisons of phenomenal character across times and across subjects.<sup>9</sup> If we can, then we can make sense of talk about introspectable differences between experiences of different subjects (or non-consecutive experiences of a single subject) in terms of introspectable differences between consecutive experiences of a single subject: there is an introspectable difference between two arbitrary experiences  $x$  and  $y$  if and only if were a subject to have consecutive experiences with (respectively) the same phenomenal characters as  $x$  and  $y$ , there would be an introspectable difference between those experiences. Because I think that this does make sense, I am inclined to think that the above argument supports all three versions of intentionalism stated above.

The arguments which follow will not, however, presuppose this. In what follows, intentionalism will figure mainly as a thesis with which our view of perceptual content ought to be consistent. For these purposes, I will be concerned with the strongest of the three theses above — interpersonal intentionalism — and will use ‘minimal intentionalism’ as a label for this view. If it holds for the interpersonal version, the consistency result will hold for the intrapersonal versions of the thesis as well.

## 2 From transparency to Russellianism

Even if we accept the argument from transparency to minimal intentionalism, Moore’s transparency point might, for all that, seem to leave open the question of what sorts of things the contents of experience are. Next I want to argue that this is not so: the transparency of experience does not just show that content fixes phenomenology, but also reveals the nature of the contents of perception.

In particular, I think that the transparency of experience provides a way of deciding between the following opposed views of content:

*Russellianism*: contents are structured objects the constituents of which are worldly items such as objects and properties.<sup>10</sup>

*Fregeanism*: contents are structured objects the constituents of which are ways of thinking about, or modes of presentation of, objects and properties.

There is wide, though not universal, agreement that these are the two most plausible views of the contents of mental states. The main alternative to this pair of views is the view that propositions are sets of possible worlds. I set this view aside for the familiar reason that there are differences in content which correspond to no differences in truth-conditions, and hence which can’t be modeled by views which identify propositions with sets of possible worlds.<sup>11</sup>

<sup>9</sup>For some skepticism on this score, see Stalnaker (1999).

<sup>10</sup>For simplicity, I am ignoring the distinction between views which take the contents of perceptions to be characterized wholly by Russellian propositions, and views which take the contents of perceptions to be characterized by such propositions in conjunction with proposition-like contents which have as their constituents the same kinds of things which are constituents of Russellian propositions. For a defense of a view of this kind, see the discussion of scenario contents and protopositions in §§3.1-3.3 of Peacocke (1992).

<sup>11</sup>As is well-known, this comes out most clearly in the identification of all necessarily true and necessarily false propositions. Examples of perceptions with necessarily true and necessarily false contents are harder to come by than examples of sentences or beliefs with these contents. Possible examples include perceptual representation of the essential properties of things (an experience which represents my desk as wooden,

The core difference between Russellian and Fregean views of content is that the Russellian thinks of contents as built up from worldly items like objects and properties, whereas the Fregean thinks of contents as built up from modes of presentation of, or concepts of, these items. There is thus a difference in kind between the sorts of things the Russellian and the Fregean take contents to be. Fregeanism is not just the denial of Millianism, understood as the view that external particulars are sometimes among the constituents of propositions; Millianism is one species of Russellianism, but not the only one. A Russellian might think, for example, that propositions are built up solely of properties, with names analyzed in terms of quantifier phrases whose contents are themselves built up out of properties.<sup>12</sup> This sort of view is anti-Millian, but not Fregean, since it makes no use of ‘modes of presentation’ or anything which answers to traditional Fregean claims about content. The argument which follows is directed against Fregeanism in the strict sense, not against anti-Millianism.<sup>13</sup>

Unlike the Russellian, who constructs propositions out of objects and properties, the Fregean introduces a new class of items — senses — to play the role of contents. It is therefore reasonable to think that the Fregean should be able to provide some constraints on when two experiences have the same Fregean sense, and when they do not. When the bearers of content in question are sentences, such constraint is standardly provided by some version of Frege’s Criterion:

Two sentences  $S$  and  $S'$  have the same sense  $\equiv$  any rational agent who understood both would, on reflection, judge that  $S$  is true just in case he would judge that  $S'$  is true.<sup>14</sup>

But how are we to apply this to the case of perception? A straightforward generalization would yield something like this:

Two experiences  $e$  and  $e'$  have the same sense  $\equiv$  any rational agent who had both experiences would, on reflection, judge that  $e$  is veridical just in case he would judge that  $e'$  is veridical.

But the latter criterion presents problems of interpretation which the former does not. Implicit in the criterion for sameness and difference of the senses of sentences is the require-

---

an experience which represents the light as metal) and the waterfall illusion, in which the subject claims that her experience represents something as both moving and not moving (Crane (1988)).

Closely related to possible-worlds views of contents are views which take propositions to be structured intensions. These views are unaffected by the above examples, since they can distinguish between truth-conditionally equivalent contents. For the purposes of the argument which follows, views of contents as structured intensions can be thought of as versions of Russellianism which identify properties with sets of objects, or functions from worlds to sets of objects.

<sup>12</sup>This view has obvious affinities to Russell’s own.

<sup>13</sup>This is important because some views of the content of experience which are called ‘Fregean’ can be understood either as Fregean in the strict sense or as non-Millian versions of Russellianism. It seems to me that the positive claims about content in Chalmers (2004), for example, could be made to fit either mold.

<sup>14</sup>An immediate problem for this statement of the Criterion is the possibility of two intuitively non-synonymous sentences which are such that anyone, on reflection, would judge them both to be true (or false). ‘ $1=1$ ’ and ‘ $2=2$ ’ do not have the same sense, even if we can’t imagine a situation in which a reflective agent would differ over their truth value. The most natural way around this problem is to allow the relevant truth-value judgements to be not only about  $S$  and  $S'$ , but also about any pair of complex sentences which differ only in the substitution of  $S$  for  $S'$ . The Criterion can naturally be extended in an analogous way to subsentential expressions, so that two expressions  $x$  and  $y$  have the same sense  $\equiv$  any rational agent who understood both would, on reflection, make the same truth value judgement about any two sentences which differ only in the substitution of  $x$  for  $y$ .

ment that the agent consider the two sentences at the same time; we can't demonstrate that 'Bob is a bachelor' differs in sense from 'Bob is an unmarried man' by noting that a friend of Bob's might rationally judge the first to be true the morning of the wedding while judging the second to be false that evening. The problem is that one can't have two experiences at the same time in the same way in which one can consider the truth of two sentences at the same time, so we can't simply build the 'same time' requirement into the second criterion. But, clearly, we need something like the 'same time' requirement in the case of the second criterion just as much as in the case of the first. After all, suppose that an agent has experience  $e$  after being told by a reliable authority that his next experience will be veridical, and experience  $e'$  just after being told by the same authority that his experience will be illusory. The fact that such an agent might well be rational to judge that the first experience is veridical and that the second is illusory hardly suffices to show that the two experiences differ in sense — any more than the example above suffices to show that 'bachelor' and 'unmarried man' do.

It seems to me that we can capture the intuitive idea behind the second criterion like this: consider any agent, complete with a set of beliefs about his environment. Now consider two possible courses of that agent's immediate future: that he has experience  $e$ , and that he has experience  $e'$ . If it could ever be the case that the agent would be rational to judge one of the experiences veridical and rational for him to judge the other not to be veridical, then the two experiences differ in sense; if it couldn't be the case that any such agent's rational judgements would differ in this way, then the two experiences have the same sense. This seems to do the work which the 'same time' requirement does in the case of the first version of Frege's criterion without requiring, incoherently, that the agent simultaneously have two experiences of the scene before him.

If we adopt this interpretation of Frege's criterion as applied to perceptual experiences, then it seems clear that any two experiences which have exactly the same phenomenal character must also have the same sense. For if two experiences have the same phenomenal character, then they would be indistinguishable to any subject. But then how could the pair of experiences supply a rational ground for any difference in veridicality judgements? They couldn't; holding fixed all relevant aspects of an agent's psychology, it can never be rational to judge one experience veridical if it would be rational, in just the same circumstances, to judge illusory an experience in principle indistinguishable from the first.<sup>15</sup> So the Fregean about the contents of perceptual experience is committed to the following:

Phenomenology/Sense Principle

If two experiences have the same phenomenology, then they have the same sense.

This commitment should not be surprising. The Phenomenology/Sense Principle is not just derivable from the most plausible way of applying Frege's Criterion to the perceptual experience; it also answers to the intuitive idea of a Fregean sense as a mode of presentation of a reference. It seems plausible that two experiences encode the same 'mode of presentation' of an object, for example the planet Venus, if there is no phenomenological difference between their sensory presentations of the planet. So visual experiences of

<sup>15</sup>Obviously, counterexamples to the idea that veridicality judgements can't rationally diverge without a divergence in phenomenology can be generated by allowing variance in the agent's circumstances, beliefs about her environment, etc. But, equally obviously, this is irrelevant to determining what follows from any plausible interpretation of the second version of Frege's criterion above.

Venus in the morning (in this part of the sky) and in the evening (in that part of the sky) might agree in reference (inasmuch as they are both experiences of Venus) while differing in their mode of presentation, or way of presenting, that object, while two visual experiences of Venus would agree in mode of presentation if there is no phenomenological difference between their presentations of the planet.<sup>16</sup>

An important consequence of the difference between Russellian and Fregean views of content is that, if the content of experience is Fregean rather than Russellian, then, since there will be many Fregean senses corresponding to each visually represented property, there will be, for each Russellian proposition attributing a property to an object, many Fregean propositions which are 'about' the same object and property, but differ with respect to the mode of presentation of the property. (There will also be different Fregean propositions which differ with respect to the mode of presentation of the object, but ignore that for simplicity.) Consider two such Fregean propositions,  $fp_1$  and  $fp_2$ , which correspond, in the above sense, to a Russellian proposition  $rp$ . The key question is then: Is the phenomenology of an experience which has  $fp_1$  as its sense different from one which has  $fp_2$  as its sense, or not?

It seems that  $fp_1$  and  $fp_2$  must differ in phenomenology. By hypothesis, they are distinct Fregean senses; and the contrapositive of the Phenomenology/Sense Principle says that difference in sense guarantees difference in phenomenology. So an experience with  $fp_1$  as content differs phenomenally from one which has  $fp_2$  as content; and this means that there is an introspectable difference between an experience which has  $fp_1$  as content and one which has  $fp_2$  as content. But now recall the moral of our discussion of the transparency of experience. Since introspection reveals only the objects that are presented as being in one's environment, and the properties those objects are presented as having, any introspectable difference between two experiences must involve some difference in the objects and properties presented as in one's environment. This is the Transparency/Difference Principle, discussed above:

Transparency/Difference Principle

If there is an introspectable difference between two experiences, then there is a difference in the objects and properties those two experiences represent as in one's environment.

It follows from this principle, along with our conclusion that there is an introspectable difference between experiences which have  $fp_1$  and  $fp_2$  as contents, that two experiences which, respectively, have  $fp_1$  and  $fp_2$  as contents differ with respect to which objects and properties they represent as before the perceiver. But this contradicts our initial supposition that  $fp_1$  and  $fp_2$  correspond to the same Russellian proposition because, if they did, they would not differ with respect to which objects and properties represented as being in the agent's environment, but only with respect to the modes of presentation of the same objects and properties. So the hypothesis that the contents of perceptions

---

<sup>16</sup>There is a subtlety here which I am ignoring for simplicity, and which plays no role in the argument which follows. One might worry that the Phenomenology/Sense Principle conflicts with the view that sense determines reference. After all, can't two experiences be phenomenally indistinguishable but of distinct objects? The Fregean has a few options here. Perhaps the most promising is to hold, following Evans, that the relevant notion of mode of presentation or way of thinking about an object guarantees sameness of reference; if there is no one object that two people are perceiving, then there is no object that the two are perceiving in the same way, and so their experiences cannot have the same Fregean sense. See the discussion of thinking about places in Evans (1982), p. 21.



are Fregean senses, along with the transparency of experience, entails a contradiction.<sup>17</sup>

It seems to me that the only plausible reply for the Fregean who wants to respect the view of the transparency of experience sketched above is to try to block the derivation of the Phenomenology/Sense Principle. One initially plausible way to do this is to relax the criterion for sameness and difference of the senses of perceptual experiences to simply a criterion of difference. On this view, differences in the relevant kinds of veridicality judgements are sufficient for difference in sense, but sameness of the relevant kinds of veridicality judgements is not sufficient for sameness of sense. So the view in question would reject the principle above, namely

Two experiences  $e$  and  $e'$  have the same sense  $\equiv$  any rational agent who had both experiences would, on reflection, judge that  $e$  is veridical just in case he would judge that  $e'$  is veridical.

in favor of

Two experiences  $e$  and  $e'$  have the same sense  $\rightarrow$  any rational agent who had both experiences would, on reflection, judge that  $e$  is veridical just in case he would judge that  $e'$  is veridical.

From the latter, there is no ready derivation of the Phenomenology/Sense Principle.

But this saves the Fregean view of experience only by making the notion of the Fregean sense of an experience unacceptably obscure. To see this, consider any two experiences which (according to the Fregean) have distinct senses as their contents, but correspond in the above sense to a single Russellian proposition. To borrow a phrase from Michael Nelson, call any pair of experiences related in this way a ‘puzzling pair.’<sup>18</sup> We know that, to block the conflict with the transparency of experience, the Fregean must claim that puzzling pairs never differ in phenomenology. (This was the point of the retreat from giving necessary and sufficient conditions for sameness of sense to giving necessary conditions.) We know from the definition of puzzling pairs that they do not differ with respect to which objects and properties they represent as being before the perceiver. But, despite the fact that puzzling pairs present just the same objects and properties in phenomenally identical ways, the Fregean claims that there is some difference between the members of puzzling pairs in their mode of presentation of some object of property. But in what could such a difference of mode of presentation consist?

It is important to be clear that the worry here is not just the standard worry that Fregeans do not say enough about the nature of the senses invoked to explain linguistic and cognitive phenomena. The worry is that, even if we grant the distinctions between the senses of expressions which Fregeans typically accept and Russellians reject, this still gives us no purchase on the differences in sense claimed to obtain between puzzling pairs of experiences.

---

<sup>17</sup>One response to this problem that I have heard is that although in general the relationship between Fregean and Russellian propositions is many-one, this may not be so in the case of the contents of perception. So, for any Russellian proposition, there is at most one Fregean proposition which is such that it is a possible content of perception and is ‘about’ the same objects and properties as the Russellian proposition. For this to be a plausible reply, we’d need some motivation for making this restriction in the case of perception, but not thought or language. One might attempt to provide such a motivation via the claim that the senses which are the contents of perceptual experiences are always demonstrative senses (as in the view of Brewer (1999)). But this is not enough; the relationship of *demonstrative senses* to any given object or property is many-one.

<sup>18</sup>The phrase is due to Nelson (2002), though Nelson has in mind certain pairs of sentences which (allegedly) differ in Fregean sense but not Russellian content, rather than pairs of experiences.

One way to show this is to compare the Fregean's claim here with standard Fregean claims in the philosophy of language. There is no shortage of pairs of sentences which the Fregean claims to differ in sense, but which the Russellian regards as having the same content; the easiest and most well-known examples are sentences which differ only in the substitution of simple coreferential names. The difference which the Fregean claims to hold between puzzling pairs of experiences is analogous to the difference which the Fregean claims to hold between such a pair of sentences: they differ in sense, but correspond to the same Russellian proposition. The Russellian, of course, will deny that there typically are such differences in sense; but, for purposes of argument, let's grant the Fregean's claim about linguistic expressions. This gives us a test for the Fregean's claim about the contents of experience.

Often, we can use language to report how a perceptual experience represents the world to us. If, as we are supposing, there are many sentences which differ in Fregean sense but not Russellian content then, if puzzling pairs of experiences really exist, one would expect that there is a puzzling pair of experiences  $e_1$  and  $e_2$  such that there is some sentence which correctly (even if partially) expresses the way that  $e_1$  represents the world as being, but does *not* correctly express the way that  $e_2$  represents the world as being.<sup>19</sup> But now consider a puzzling pair of experiences, which represent just the same objects and properties as before the perceiver, and are phenomenally identical. Is it ever the case that a sentence would be a correct report of the content of one member of that pair of experiences, but not of the other? It seems to me that the answer is 'No'.

A similar argument can be run at the level of thought. The Fregean will hold that the contents of thoughts are often distinct, even though identical at the level of Russellian content. So, if there are puzzling pairs of experiences, we should expect that sometimes a thought would (even if partially) match in content one member of the pair, but not the other. But again consider a puzzling pair of experiences, which represent just the same objects and properties as before the perceiver, and are phenomenally identical. Is it ever the case that a thought or judgement would represent the world in the same way as one of those experiences, but not the other? If, as I think, the answer is again 'No', the Fregean is stuck with the following view about puzzling pairs of experiences:

There are indefinitely many pairs of experiences related in the following way: they present just the same objects and properties in phenomenally identical ways, but nonetheless differ in their mode of presentation of those objects and properties. Sentences and thoughts can also differ in their modes of presentations of objects and properties, but the differences in mode of presentation we find in the case of perception are undetectable because they are both inaccessible to thought and inexpressible in language.

This view does not conflict with the transparency of experience; but that seems to be all that it has to recommend it.

So the differences in content posited by the Fregean conception of content either, if accompanied by differences in phenomenology, conflict with the transparency of experience, or, if unaccompanied by differences in phenomenology, make the idea of a Fregean

---

<sup>19</sup>Otherwise, the Fregean would have to claim that, although there are differences in the senses of experiences which are not reflected as differences in Russellian content, and there are differences in the senses of linguistic expressions which are not reflected as differences in Russellian content, of necessity the two never coincide. But what could explain that?

sense unintelligible. The best conclusion is that the differences in content posited by the Fregean conception of the content of experience do not exist.<sup>20</sup>

### 3 *Is Russellianism inconsistent with intentionalism?*

The argument so far might seem to lead less to a satisfactory view of the contents of experience than to a kind of antinomy. Reflection on the transparency of experience leads directly to minimal intentionalism and, by an only slightly longer route, to a Russellian view of the contents of experience. But it might seem that these two views, rather than being natural companions, are inconsistent.

Intentionalism suggests that there is a close link between the content and phenomenology of a perceptual experience. But, if we examine a simple Russellian proposition — say one which attributes a color to an object — it is easy, given the existence of perceptually indistinguishable objects, to find cases in which two experiences differ in content but not in phenomenology. So, it might seem, either Russellianism is false or intentionalism is, and the argument from the transparency of experience to these two conclusions must have gone wrong somewhere.<sup>21</sup>

There is a quick response to this argument. Above I argued that the transparency of experience provides a good argument for minimal intentionalism, which I characterized as the view that any difference in phenomenology must correspond to a difference in content, i.e.:

#### Minimal intentionalism

If two experiences differ in phenomenal character, then they differ in content.

But the argument above does not show that *this* view is inconsistent with a Russellian view of content. Minimal intentionalism rules out changes in phenomenology which are not accompanied by changes in content; but it does not rule out changes in content which are not accompanied by changes in phenomenology. The above argument does indicate that Russellian views of content are incompatible with a stronger form of intentionalism:

#### Biconditional intentionalism

Two experiences differ in phenomenal character if and only if they differ in content.

But the transparency of experience does not provide an argument in favor of biconditional intentionalism; so the argument from transparency to the conjunction of minimal intentionalism and the Russellian view of content appears safe.

But this response to the argument is a bit *too* quick. It is convincing only if minimal intentionalism without biconditional intentionalism is a stable position; and many have thought that it is not.

One way of developing this worry is via an argument that any motivation for being a minimal intentionalist also supports biconditional intentionalism. One guiding thought

---

<sup>20</sup>Thanks for helpful discussion of Fregean views of content to Ben Caplan, Dorothy Edgington, Ben Morison, and Richard Price.

<sup>21</sup>One might try to respond to this argument by denying that objects can be parts of the contents of experience, and retreating to what I above called an anti-Millian version of Russellianism. However, this only delays the problem. For a convincing argument to this effect, see Caplan and Schroeder (forthcoming).

behind intentionalism is this: though the phenomenal character and content of an experience are conceptually distinct, it would be odd if there were *no* systematic connection between these aspects of perception; it would be odd if there was nothing to be said about the connection between how an experience seems to a perceiver and how that experience represents the world as being. But if there is some systematic connection between the contents and phenomenal characters of perceptual experiences, it seems plausible that one of these two aspects of experience should explain the other. Many theorists think of the contents of mental states as being more theoretically tractable than the phenomenal character of perceptions, and so have wanted to explain facts about the phenomenal character of experience in terms of the contents of experience.

So far, this thought might not seem to pose any special challenge to minimal intentionalism. Minimal intentionalism says that any difference in phenomenology corresponds to a difference in content; so it is natural for the minimal intentionalist to adopt the principle that differences in phenomenology are always to be explained by differences in content. But then we might worry that

“Holding this principle seems to require holding the corresponding principle about sameness. (Otherwise, differences in content would sometimes be used to explain sameness of qualitative character and sometimes be used to explain differences in qualitative character.)”<sup>22</sup>

This argument poses a dilemma for the minimal intentionalist position. Given that differences in phenomenology are to be explained in terms of differences in content, how are we to explain sameness in phenomenology? If we explain sameness in phenomenology in terms of sameness of content, then we are led from minimal intentionalism to biconditional intentionalism. But if we say that sameness in phenomenology is consistent with difference in content, then difference in content is consistent with both sameness and difference of phenomenal character, which seems to compromise the wanted explanation of difference of phenomenology in terms of difference of content.

This argument is not convincing. Consider the following view about dalmatians:

Every dalmatian has some number of spots. If you change the number of spots a dalmatian has — by dyeing some of its hair black, for example — you have to change the color of at least one of its hairs. So every change in the number of spots brings with it a change in the color of one or more hairs; but the converse doesn’t hold, since you could change the color of one hair of a dalmatian without adding to its spots.

This view about dalmatians is the analogue of minimal intentionalism without biconditional intentionalism: it holds that a change in the  $x$ ’s brings with it a change in the  $y$ ’s, but that a change in the  $y$ ’s is consistent with sameness or difference of  $x$ ’s. But this, fairly obviously, does not preclude explaining changes in the  $x$ ’s — the number of spots a dalmatian has — in terms of changes in the  $y$ ’s — the color of the hairs of the dalmatian. This illustrates one way in which the above argument against minimal intentionalism might fail: content might be related to phenomenology as color of dalmatian hairs is related to numbers of dalmatian spots. It might be that a certain degree of change in content brings about and explains a change in phenomenology, but that small changes in content do not always bring about changes in the phenomenology.

---

<sup>22</sup>Caplan and Schroeder (forthcoming).

A possible example of this sort of relationship comes from the content and phenomenology of experiences of the colors of things. Many have thought that we can construct a ‘phenomenal continuum’ of color patches which are such that, although each adjacent pair is indistinguishable in respect of color, the first clearly differs in color from the last. If there were such a series, it would fit the model above. Since an experience of the first patch differs from an experience of the last in what color it represents the patches as being, it seems that there must be at least two adjacent patches  $x$  and  $y$  which are such that an experience of  $x$  would represent it as having a different color than an experience of  $y$  would represent it as having. But, by hypothesis, experiences of the colors of  $x$  and  $y$  are subjectively indistinguishable. So here we have an example of two experiences which differ in content, but where the difference in content is in an intuitive sense not enough to generate a difference in phenomenology.<sup>23</sup>

But there is another way in which the argument might fail. Consider the following view about baseball games:

The actions of umpires determine how many outs there are in an inning; so, with every change in the number of outs, there must also be a change in the actions of the umpire. But the converse does not hold; the umpire can perform an action — calling a pitch a ball, for example — without changing the number of outs in the inning.

As above, this asymmetric supervenience of outs on umpire actions hardly precludes explaining the number of outs in an inning in terms of the actions of the umpire. But the relationship between umpire actions and the number of outs is different than that between the colors of dalmatian hairs and the number of spots a dalmatian has; it’s not that ‘small enough’ changes in umpire actions don’t yield a change in the number of outs, but rather simply that some kinds of actions of the umpire don’t affect the number of outs, while others do.

This provides another model of the relationship between content and phenomenology which is consistent both with minimal intentionalism and with the view that differences in phenomenology are to be explained in terms of differences in content: perhaps the contents which can figure in the contents of perceptual experience come in two sorts: phenomenology-affecting and phenomenology-silent. A kind of content would be phenomenology-silent if any two experiences which differ *only* with respect to that kind of content cannot differ in phenomenology. Then changes in phenomenology could be explained by changes in the phenomenology-affecting content of an experience, while sameness of phenomenology would be consistent with changes in the phenomenology-silent aspects of the content of an experience.

Examples of pairs of experiences in which content and phenomenology are related in this way are not hard to come by if, as seems plausible, the contents of experience can include external particulars. External particulars seem to be, in the above sense, phenomenology-silent. For consider two experiences which differ only in the objects they are experiences of. Since these experiences do not differ with respect to the properties they represent these objects as having, these objects will, for the purposes of the experiences in

---

<sup>23</sup>The argument is generated by the fact that even though indiscriminability seems not to be transitive, it is hard to see how sameness of content could fail to be transitive. For an argument that indiscriminability is transitive, see Graff Fara (2001). If indiscriminability is transitive, then the above argument fails; in that case it is not obvious that there are any contents which are related to phenomenology in the way illustrated by the example of the dalmatian.

question, be indistinguishable. So the two experiences will not differ in phenomenology, even though they differ in content, because their difference in content is not of the right sort to register at the level of phenomenology.<sup>24</sup>

Conclusion: even if the minimal intentionalist does aim to explain differences in phenomenology via differences in content, this explanatory ambition does not motivate the move from minimal intentionalism to biconditional intentionalism.

But there is a different kind of motivation for intentionalism which might seem to put more pressure on the minimal intentionalist. Intentionalism is sometimes presented as the last, best hope of materialism about the mind. The argument goes something like this: our intuitions about phenomenal character — in particular, intuitions about the possibility of zombies and spectrum inverted subjects who are physically identical to normal perceivers — make it very hard to see how we could give a materialist account of phenomenal character. But we do have a reasonably good idea how to give a materialistically acceptable theory of mental content. For an internal state to have a given content is just for . . . (fill in your favorite theory of mental content here). So for the purposes of defending materialism, it would be very convenient if we could just *reduce* phenomenal character to content. To this end, perhaps we should hold that properties of experiences like ‘having such-and-such phenomenal character’ are just identical to properties of experiences like ‘having such-and-such content.’ As Michael Tye puts it, “phenomenal character is one and the same as representational content that meets certain further conditions.”<sup>25</sup> Then if we can give a materialistically acceptable account of the latter, we will have succeeded in giving a materialistically acceptable account of the former.

The problem, for present purposes, is that identity is symmetric. If the property (of experiences) of having a given phenomenal character really is identical to the property of having a given content, then any change in one entails a change in the other. And that is biconditional intentionalism, which is what we have been trying to avoid.

One might simply abandon materialism; one might also doubt whether we really do have reason to be confident that a materialistically acceptable theory of mental content is in the offing. But there are two more fundamental worries about this case for biconditional intentionalism.

First, there is something odd about the idea that a property identity of the relevant sort might help the case of materialism. Suppose that content were, in some sense, identical to phenomenal character. This would not make our intuitions about zombies and spectrum inversion simply go away. If it is extremely difficult to give an account of phenomenal character, this task does not automatically get easier if we find that phenomenal character = representational content. If anything, the truth of such a property identity would seem to show that giving a physicalist account of representational content is much harder than we thought. So there is reason to distrust at least one of the motivations behind this reductive program.<sup>26</sup>

Second, the denial of biconditional intentionalism does not rule out the idea that

---

<sup>24</sup>Philosophers reluctant to countenance changes in content which are not accompanied by changes in phenomenology might be tempted to deny, on this ground, that external particulars can figure in the contents of perceptions. I do not think that this is plausible; I return to this point below.

The view that perceptual representation of external particulars is phenomenology-silent does face at least *prima facie* difficulty accounting for cases of prosopagnosia. I’m not sure what the right thing to say about this case is.

<sup>25</sup>Tye (2002), p. 137. For a clear statement of this motivation for the identification of content and phenomenal character, see Lycan (1996).

<sup>26</sup>For a defense of the contrary view that such a property identity could help to solve the problems posed by intuitions about zombies and other duplicates, see §7.4 of Tye (1995).

there might be property identities which would reduce facts about phenomenal character to some other class of facts. One idea here is that we could explain phenomenal character in terms of the objects, rather than the contents, of experience; another is that we could explain phenomenal character in terms of the phenomenology-affecting aspects of the contents of experience. Either reductive strategy would be consistent with the truth of minimal intentionalism, and the falsity of biconditional intentionalism.

#### 4 *Against biconditional intentionalism*

So far I have argued that reflection on the transparency of experience gets us to a conjunction of views: minimal intentionalism and Russellianism about the contents of experience. I have been suggesting that we can avoid problems with this conjunction by resisting the move from minimal intentionalism to biconditional intentionalism. I now want to argue that there are independent reasons for rejecting biconditional intentionalism.

Biconditional intentionalism is the conjunction of minimal intentionalism with the following thesis:

If two experiences differ in content, then they differ in phenomenal character.<sup>27</sup>

The problem with this principle is that it makes it hard to see how perceptual experience could play a role in our cognitive lives which, as Mark Johnston has emphasized, it clearly does play: making new contents available for thought.<sup>28</sup>

Whatever one's view of content, there is some distinction to be captured between thinking about an object *de re*, and merely thinking about it under a description. Given this distinction, it seems clear that there are some external particulars about which I am in a position to entertain and express *de re* thoughts, and that there are some external particulars with respect to which I am not in such a position. It also seems clear that an object can become available for *de re* thought by my having a perceptual experience of that object. Before moving to Montreal, I was not in a position to have *de re* thoughts about the Jacques Cartier bridge — but seeing the bridge changed that.<sup>29</sup>

Furthermore, it seems clear that which objects an experience makes available for *de re* thought should have something to do with the content of the experience. It is intuitively plausible that an experience which represents an object as being a certain way should make available *de re* thoughts about that object. Different views of the *de re* thought and different views of the nature of the contents of perceptions and of thoughts will characterize this connection between perception and availability for thought differently; but the following core idea is difficult to deny:

##### Perception/Availability Principle

If two experiences differ in which thoughts they make available to the subject of the perception, then they differ in content.

For suppose that two experiences have just the same content. Then they represent the world as being just the same way; they present just the same objects as having just the

---

<sup>27</sup>This principle is not the same as, but is clearly related to, the 'Phenomenal Bottleneck Principle' argued against in §11 of Johnston (2004).

<sup>28</sup>See especially Johnston (2004, 2006).

<sup>29</sup>I'm not suggesting that perception is the *only* way for something to become available for *de re* thought — just that it is one way.

same properties. How could one make available thoughts about an object or property which the other did not?

If the Perception/Availability Principle is true, then biconditional intentionalism is false. For there can be distinct but perceptually indistinguishable objects, and a perceptual experience could put one in a position to have de re thoughts about one but not about the other, from which it follows that an experience of one can differ in content from an experience of the other. But if the objects really are perceptually indistinguishable, these two experiences will not (or need not) differ in phenomenology.<sup>30</sup>

The defender of biconditional intentionalism can reply by saying that the availability of an object for de re thought is sometimes explained by an aspect of a perceptual experience other than its content. She might say, for example, that an object can become available for de re thought by being involved in the causation of a perceptual experience, even if the experience does not represent that object. But, apart from the ‘extensional’ problem of delimiting this causal condition so that it provides a plausible sufficient as well as necessary condition on de re acquaintance,<sup>31</sup> it is hard to see how causing an experience without affecting its content could matter for de re acquaintance. Moreover, the suggestion looks like the introduction of an epicycle. If we take seriously the idea that perceptual experiences have contents, then we should want to avoid explaining how perceptions make contents available for thought partly in terms of their contents, and partly in terms of their non-content-related causes. Biconditional intentionalism thus comes with a significant cost: it severs a plausible connection between perception and availability for thought.

## 5 Externalism about content and internalism about phenomenology

A related form of argument can be used to show that the contents of perceptual experiences are externalist — that two agents can be alike in all their intrinsic qualities, and yet be having experiences which differ in content. Consider again two experiences, which differ only in that one is of an object  $x$ , and the other is of an indistinguishable object  $y$ . The first puts the subject of perception in a position to have de re thoughts about  $x$ , but not  $y$ ; the second puts the subject of the perception in a position to have de re thoughts about  $y$ , but not  $x$ . So, by the Perception/Availability Principle, the two experiences differ in content. But, given that  $x$  and  $y$  are indistinguishable, it is clear that this difference in content need not show up as a difference in the intrinsic properties of the subjects of the two experiences. So the contents of perceptions do not supervene on the intrinsic properties of the subjects of those perceptions.<sup>32</sup>

But externalism about the contents of perceptions can seem to pose a problem for the intentionalist view which I have been defending. For intentionalism claims that there is a systematic connection between the content and phenomenal character of experience. But any attempt to draw a connection between these two aspects of experience faces a problem in making sense of how, if the contents of perceptual experience are determined by factors external to the subject, the phenomenal character of a perceptual experience can supervene on the intrinsic properties of the perceiver.

<sup>30</sup>Substantially the same point could be made using as examples perceptually indistinguishable natural kinds, rather than perceptually indistinguishable particulars.

<sup>31</sup>These are related to the problems with developing a causal theory of perception on the model of Grice (1961/1989); for a nice summary, see the opening section of Noë (2003).

<sup>32</sup>It is worth noting that this argument for externalism does not rely on Russellianism, or any contentious views about content. It relies only on the idea that there is some distinction between de re and descriptive thought, and on the Perception/Availability Principle.



The intuition behind the view that the phenomenal character of an experience is fixed by the intrinsic properties of the perceiver is something like this: standard externalist thought experiments show that by changing facts external to an agent, one can change the contents of his mental states and the meanings of expressions of his language. But you can't change how things seem, subjectively, to the subject himself. *That* is fixed by how the subject is, not by any facts about his physical or social environment.<sup>33</sup> This view can be expressed as follows:

Internalism about phenomenology

It is impossible for two experiences to differ in phenomenal character without a change in the intrinsic properties of the perceiver(s).

There is a tension between this principle, intentionalism, and externalism about content: if, as intentionalism says, there's a necessary connection between phenomenology and content, how can one of these supervene on intrinsic properties, and the other not?

It is easy to turn this intuitive tension into an argument against biconditional intentionalism. Consider two intrinsic duplicates. By externalism about the content of perception, their perceptions might have different contents. Then, by biconditional intentionalism, their experiences have different phenomenal characters. But this contradicts internalism about phenomenal character.

But this argument should not impress the minimal intentionalist, who will reject the inference from a difference in content to a difference in phenomenology. How could this argument be modified to work against a minimal intentionalist?

One plausible route, taken by Andy Egan and James John in "A Puzzle About Perception", is to strengthen the externalist thesis.<sup>34</sup> The kind of externalism suggested above was equivalent to the following:

Weak externalism

$\exists p$  (a subject could be having an experience involving content  $p$  even though his molecule-for-molecule duplicate is not)

In particular, the example above involving perceptually indistinguishable objects seems to show that the contents of experiences which provide the ability to have *de re* thoughts about external particulars — however exactly we characterize the contents of such experiences — do not supervene on the intrinsic properties of the perceiver.

It seems plausible that the truth of weak externalism can only be explained by the fact that sometimes the content of a subject's experience is shaped by facts external to him. This might suggest the following view: what it is for a subject to represent the world in a certain way just is for him to be connected to, or embedded in, his environment in a certain way. But then it might seem that we can advance a stronger claim than the bare externalist thesis above; perhaps it's not just the case that some contents of experience fail to supervene on the intrinsic properties of perceivers, but that *all* such contents fail to supervene on intrinsic properties of perceivers. We can express this stronger externalist view as follows:

---

<sup>33</sup>For some evidence of the strength of this intuition, note that in all of the original externalist thought-experiments, it was assumed that the changes in environmental or social facts which were claimed to explain changes in content would leave 'how things seem to the subject' fixed. Putnam, for example, did not hesitate to assume that it could be the case that "Oscar<sub>1</sub> and Oscar<sub>2</sub> were exact duplicates in appearance, feelings, thoughts, interior monologue, etc." (Putnam (1973), 702).

<sup>34</sup>Egan and John (ms.).

### Strong externalism

$\forall p$  (a subject could be having an experience involving content  $p$  even though his molecule-for-molecule duplicate is not)

Let's suppose that this is true. Then it seems as though two perceivers might be such that the correlations between their internal states and certain of the contents of their perceptions are inverted. So, for example, it might be that when I am in intrinsic state  $I$  I am representing  $o$  as red, but when the agent content-inverted relative to me is in state  $I$ , he is representing  $o$  as green. (It is important to see that this is different from the usual *spectrum* inversion cases, which are inversions of phenomenal character, not content.)

Supposing that content inversion is possible, we then argue as follows.<sup>35</sup> Let Nonvert be a perceiver who is normal (by our lights), and let Invert be content-inverted relative to Nonvert.<sup>36</sup> At time  $t_1$ , let Nonvert have intrinsic properties  $I$  and be having a perceptual experience which has the content that some object  $o$  is red. At time  $t_2$ , let Nonvert be having an experience which represents  $o$  as green. It then seems clear that Nonvert's experience will have different phenomenal characters at  $t_1$  and  $t_2$ , since when we have an experience which represents an object as red, it differs in phenomenal character from an experience which represents that object as green. Now suppose that Invert is having a perceptual experience, and is an intrinsic duplicate of Nonvert at  $t_1$ . Because Invert is content-inverted relative to Nonvert, his experience will represent  $o$  as green. By minimal intentionalism, Invert's experience will have the same phenomenal character as Nonvert's at  $t_2$ . But then since Nonvert's experiences at  $t_1$  and  $t_2$  have different phenomenal characters, it follows that Invert's experience has a *different* phenomenal character than Nonvert's experience at  $t_1$ . But this contradicts internalism about phenomenal character, since Invert has exactly the same intrinsic properties as Nonvert at  $t_1$ .

The conjunction of strong externalism and internalism about phenomenal character therefore entails the falsity of minimal intentionalism. But there is reason to be skeptical about this conjunction of theses.

For one thing, it places quite specific demands on our intuitions about what is and what is not possible. One way of seeing this is that it in effect assumes that content inversion is possible, but spectrum inversion is not, since were spectrum inversion (of certain sorts) possible, internalism about phenomenal character would be false. But it is hard to find a ground for supposing one possible but the other not.<sup>37</sup>

A second way of casting doubt on this conjunction is to show that it proves too much. The negation of minimal intentionalism just says that *sometimes* sameness of content can go with difference of phenomenal character; the conjunction of strong externalism and internalism about phenomenology entails that, in the case of content-inverts, sameness of content *must always* go with difference of phenomenal character. But it is hard to see why this claim is more plausible than the denial of the conjunction of strong externalism with internalism about phenomenal character. So one reasonable attitude toward this argument

---

<sup>35</sup>Again, the argument here follows Egan and John (ms.). There are clearly similarities between this argument and the well-known 'inverted earth' example of Block (1990). Alleged examples of phenomenal inversion without misrepresentation are an important challenge to the combination of views defended here, but are beyond the scope of this paper.

<sup>36</sup>The names are borrowed from Byrne and Hilbert (1997).

<sup>37</sup>One possible ground for thinking content inversion possible which would not carry over to phenomenal inversion is confidence in some externalist theory of mental content which entails the possibility of content inversion, e.g. the view that content is explained by covariation under certain conditions. I do not think that any view of this sort is likely to be true, but do not argue the point here. For a powerful presentation of the motivations for a covariational approach to content, see Stalnaker (1984).

would therefore be to deny the conjunction of strong externalism and internalism about phenomenal character, without worrying too much about which conjunct makes it false.

I suggest that, whether or not internalism about phenomenal character is true, strong externalism is false.<sup>38</sup> Recall the argument given above for (weak) externalism: suppose that we take two indistinguishable objects  $x$  and  $y$ . An experience of  $x$  must have a different content than an experience of  $y$ , since such an experience would put you in a position to have *de re* thoughts about  $x$  but not about  $y$  — and it is hard to see how two experiences with just the same content, which represent the world in just the same way, could make different objects available for thought. But substituting  $x$  for  $y$  as the object of the experience needn't bring about any intrinsic change in the perceiver, despite the fact that it does bring about a difference in content. It seems clear that the argument works just as well if  $x$  and  $y$  are distinct but indistinguishable natural kinds. But, crucially, there is no obvious way to run the argument if  $x$  and  $y$  are, for example, distinct color properties.

One might try to adapt the argument like this: imagine two patches which are distinct shades of red, say RED<sub>3484</sub> and RED<sub>3485</sub>, but indistinguishable to the human eye. Surely if the shades are so close together as to be indistinguishable to the eye, switching one for the other needn't correspond to any intrinsic change in the properties of the perceiver. But experiences of the two shades differ in content. After all, if we lined up a series of such distinct but indistinguishable shades so that experiences of the first and last differed in phenomenal character, then (by minimal intentionalism) experiences of the first and last would differ in content. So there must be a difference in content between two adjacent patches somewhere in the series; and if you don't think that it falls between RED<sub>3484</sub> and RED<sub>3485</sub>, then pick the adjacent shades between which it does fall, and run the argument there.

But this argument is not convincing. Consider again the experiences of the first and last patches in the continuum. These experiences differ in phenomenology. Either this difference corresponds to a difference in the intrinsic properties of the perceiver, or it does not. If it does not, then phenomenal character does not supervene on intrinsic properties, and the present argument against intentionalism is defused. If it does, then there is a difference in the intrinsic properties of the perceiver between experiences of two adjacent shades in the series. What is the argument that this difference falls in a different spot than the difference in content?

So (what I take to be) the most powerful argument in favor of externalism about the contents of perceptions does not provide any support for strong externalism. Where, one might ask, should we draw the line between contents which are 'externalist' and those which are 'internalist' — i.e., between those contents which are such that having a perceptual experience involving them does not supervene on intrinsic properties, and those which are such that having a perceptual experience involving them does so supervene? We should look to externalist arguments, and see what conclusions they support. It is very easy to generate externalist arguments in the case of perceptual representations of external particulars or natural kinds, but, as the above discussion shows, these arguments do not clearly carry over to the cases of color and shape properties. A reasonable view therefore seems to be that particulars and natural kinds are externalist, and that 'sensible qualities' like color and shape properties are internalist.

Before I suggested, using the analogy with the baseball umpire, that some contents of

---

<sup>38</sup>For discussion of externalism about phenomenal character, see Dretske (1995), Lycan (1996), Tye (1995), Byrne and Tye (forthcoming).

experience could be phenomenology-affecting, while others could be phenomenology-silent. If the above distinction between externalist and internalist contents is to do the required work in defusing the objection from internalism about phenomenal character, we should expect the externalist contents to be phenomenology-silent, and the internalist contents to be phenomenology-affecting. (Otherwise, we could run the content-inversion argument above using the externalist but phenomenology-affecting properties.) And this is precisely what we find: it is easy to imagine pairs of experiences with the same phenomenology but differing with respect to perceptually indistinguishable particulars and natural kinds — and these are just the contents which can be shown to be externalist. It is important to see that the coincidence of externalist and phenomenology-silent properties is not a matter of stipulation. We find out what the externalist contents are by seeing where we can run externalist arguments using the Perception/Availability Principle; we find out which contents are phenomenology-silent via introspection. That they coincide lends credence to the view that the distinctions are genuine.

## 6 *Narrow content?*

This division between two kinds of perceptual contents might give rise to the following objection:

This view says that the contents of perceptions come in two classes, those which supervene on intrinsic properties and are phenomenology-affecting, and those which do not supervene on intrinsic-properties and are phenomenology-silent. Call the former class ‘narrow contents.’ Since the narrow contents of perception are all phenomenology-affecting, the narrow content of perception is a level of content for which biconditional intentionalism is true. But this shows that narrow contents cannot be Russellian contents. For consider: surely the narrow contents of perception, like any contents, can be or fail to be veridical. Now consider a perception which represents as instantiated some supposedly internalist, phenomenology-affecting property like a particular shade of red. It is not sufficient for such an experience to be veridical that that shade be instantiated somewhere or other; it is veridical if and only if the property is instantiated at the time of the perception and in the location of the perceiver. So, if Russellianism is true, these times and locations (rather than, as the Fregean would have it, modes of presentations of those times and locations) must be part of the contents of the relevant perceptions. But times and locations present the same problems for the biconditional intentionalist as do external particulars — distinct times and places can be qualitatively indistinguishable. So either a view which posits a distinction between externalist and internalist contents is incorrect, or there is a level of perceptual content — the most fundamental level — which is not Russellian.<sup>39</sup>

There are two reasons why this argument should be resisted.

The less important of these is that even if there is a class of propositions every constituent of which is phenomenology-affecting, it does not follow that biconditional intentionalism would be true of perceptions with those propositions as contents. Recall from §3 that there are two ways in which changes in content might fail to show up at the level

---

<sup>39</sup>This follows an argument which is convincingly presented in Caplan and Schroeder (forthcoming), §§3-6, though their target is not the view described in the preceding section.

of phenomenology. One is that the change is a change solely in phenomenology-silent contents; but the other is that the change in content is too slight to be evident to introspection. If this is possible — and it is hard to see what reason one could have for being sure that it is not — then even if the contents of perception were all internalist and phenomenology-affecting, biconditional intentionalism would be false.<sup>40</sup>

But the more important reason for resisting the above argument is that the move from the view that some contents are internalist to the view that there is a level of ‘narrow content’ is far from innocent. Suppose that color properties are internalist contents, and let RED be some determinate shade of red. Let

$$\langle x, \text{RED}, y, z \rangle$$

be the content of the experience of some agent *A*; suppose that *x*, *y*, and *z* are, unlike RED, externalist contents. Perhaps they are, respectively, an external particular, a location, and a time. Then this proposition is *mixed*, in the sense that it contains both internalist and externalist contents.

What follows from the supposition that RED is internalist? It follows that any intrinsic duplicate of *A* would be having an experience with a content of the form

$$\langle \_, \text{RED}, \_, \_ \rangle$$

where the blanks are filled in with an object, a location, and a time. It does not, of course, follow that this object, location, and time should be *x*, *y*, and *z*; these contents are externalist, and hence not guaranteed to appear in the contents of the perceptions of intrinsic duplicates of *A*. And, more importantly for present purposes, it does not follow from the fact that RED is internalist that there should be some *pure* proposition of which RED is a part which is a content of the agent’s experience and contains no admixture of externalist contents. In other words: it does not follow from the fact that some contents are internalist that there is a level of narrow content.

## 7 Content and hallucination

So far I have argued that the transparency of experience shows that the contents of perceptual experiences determine their phenomenology, and that the content of experience is Russellian. In replying to two challenges to this conjunction of views — the pressure to move from minimal intentionalism to biconditional intentionalism, and the argument from externalism about content and internalism about phenomenology — I have introduced two distinctions: the distinction between phenomenology-silent and phenomenology-affecting contents, and the distinction between externalist and internalist contents. In order to do the work for which they were introduced, these two distinctions would have to coincide; and I have argued that they do. The fact that they coincide lends some credence to the reality of the distinctions. But one might still object that this combination of views — Russellianism and minimal intentionalism — can only be saved by positing distinctions for which there is not enough independent evidence.

<sup>40</sup>Phenomenology-affecting contents are not defined as a class of contents which are such that any two experiences which differ in their contents with respect to members of that class *must* also differ in phenomenology; they are a class of contents which are such that two experiences which differ in their contents only with respect to members of that class *can* differ in phenomenology.

Ideally, then, we should want some grounds, independent of Russellianism and intentionalism, for positing a distinction between externalist and phenomenology-silent contents, on the one hand, and internalist and phenomenology-affecting contents, on the other. Such independent grounds are provided by a crucial asymmetry, noticed by Mark Johnston, between veridical and hallucinatory experience: there are some contents with which one can become originally acquainted in either hallucinatory or veridical experience, and others with which one can become originally acquainted only in veridical experience. In this sense, some contents are *hallucination-available* whereas others are not.

Why should this fact matter for a discussion of the nature of perceptual content and its relation to phenomenology? For our purposes, two aspects of hallucinatory experience are especially salient:

- In hallucinations, the standard external accompaniments to having a perceptual experience with a given phenomenology are absent.
- For any given veridical experience, there is a phenomenally identical hallucination.

The first aspect of hallucination should lead us to expect that, if there really is a distinction between externalist and internalist contents of perception, that internalist but not externalist contents should be hallucination-available. The second aspect of hallucination should lead us to expect that, if there really is a distinction between phenomenology-silent and phenomenology-affecting contents, that all phenomenology-affecting contents should be hallucination-available. And this, strikingly, is exactly what we find.

Sensible qualities like colors are the paradigm case of hallucination-available contents. Johnston illustrates the point with the example of super-saturated red. He imagines that

“the paradigm red – the reddest of the reds – can only be presented in delusive experience. One can come to know what supersaturated red is like only by afterimaging it. While one is afterimaging it, one could compare how much more saturated it is than the reds exhibited by the reddest of the standard Munsell color chips, there before one on the table. Likewise, a painter might discover in hallucination a strange, alluring color, which he then produces samples of by mixing paints in a novel way. Here we have all the signs of *de re* knowledge of quality.”<sup>41</sup>

Here we have a hallucinatory experience which contains a given quality — and, as a result of that experience, we are put in a position to have thoughts involving that quality. It is hard not to agree with Johnston that “I can secure my first singular reference to the quality cherry red or to the structural property C major by way of hallucinating a scene or a tune.”

But, as Johnston notices, there is a sharp contrast here with the cases of external particulars and natural kinds. He puts the point like this:

“... although we can hallucinate real things and real people, no such hallucination could be an original source of *de re* thought about those particular things or people. In this way, hallucination differs from veridical sensing, which characteristically provides new particulars as topics for thought and talk. I can hallucinate my mother talking to me on the phone, but I could not do this unless I already had an independent way of making singular reference to my

---

<sup>41</sup> Johnston (2004), p. 130.

mother. If I had been abandoned to the monks at birth and knew nothing of my mother or of mothers in general, then I could not hallucinate my mother talking to me. Even if I hallucinated a woman who happened to look just like my mother, there would be nothing that would make that hallucination of my mother, as opposed to my aunt, or any other woman who appeared like her.”<sup>42</sup>

It is not difficult to make a parallel case for natural kinds.

This strikes me as an extremely strong argument in favor of the present approach to the contents of perception. The nature of hallucination, in particular the fact that hallucinations can be had in the absence of the usual connections to the external world, makes it such that if there are externalist contents, we should expect these to be hallucination-unavailable. We have seen that we can give strong arguments that objects and natural kinds are externalist, and these are exactly the contents which turn out to be hallucination-unavailable. The nature of hallucination, in particular the fact that hallucinations can be phenomenally indistinguishable from veridical experiences, makes it such that if there are phenomenology-affecting contents, we should expect these to be hallucination-available. Colors and shapes are paradigm cases of phenomenology-affecting contents, and these turn out to be hallucination-available.

This way of stating the view poses a challenge to its rivals: could it really be an accident that when we look at our three distinctions between possible contents of perception —

internalist vs. externalist  
phenomenology-affecting vs. phenomenology-silent  
hallucination-available vs. hallucination-unavailable

— they draw the line in exactly the same place?<sup>43</sup>

---

<sup>42</sup>Johnston (2004), p. 129.

<sup>43</sup>Thanks for helpful discussion of previous versions of this paper to Casey O’Callaghan, Scott Soames, Fritz Warfield, Ben Caplan, Bob Hanna, and audiences at the University of Notre Dame, Auburn University, the University of Quebec at Montreal, and Oxford University. Special thanks also to David Davies and the other participants in my seminar on perception at McGill in the winter of 2006.

## References

- Alston, William. 2005. Perception and Representation. *Philosophy and Phenomenological Research* 70:253–289.
- Block, Ned. 1990. Inverted Earth. *Philosophical Perspectives* 4:52–79.
- Brewer, Bill. 1999. *Perception and Reason*. Oxford: Clarendon Press.
- Brewer, Bill. 2006. Perception and Content. *European Journal of Philosophy* 14(2):165–181.
- Byrne, Alex. 2001. Intentionalism Defended. *Philosophical Review* 110:2:199–240.
- Byrne, Alex and Hilbert, David. 1997. Colors and Reflectances. In *Readings on Color: Volume 1, The Philosophy of Color*, edited by Alex Byrne and David Hilbert. MIT Press.
- Byrne, Alex and Tye, Michael. forthcoming. Qualia Ain’t In the Head. *Noûs* .
- Campbell, John. 2002. *Reference and Consciousness*. New York: Oxford University Press.
- Caplan, Ben and Schroeder, Tim. forthcoming. On the Content of Experience. *Philosophy and Phenomenological Research* .
- Chalmers, David. 2004. The Representational Character of Experience. In *The Future for Philosophy*, edited by Brian Leiter. Oxford University Press.
- Crane, Tim. 1988. The Waterfall Illusion. *Analysis* 48.
- Dretske, Fred. 1995. *Naturalizing the Mind*. Cambridge, MA: MIT Press.
- Egan, Andy and John, James. ms. A Puzzle About Perception .
- Evans, Gareth. 1982. *Varieties of Reference*. Oxford: Oxford University Press.
- Fara, Delia Graff. 2001. Phenomenal Continua and the Sorites. *Mind* 110(440):905–935.
- Grice, Paul. 1961/1989. The Causal Theory of Perception. In *Studies in the Way of Words*, 224–247. Cambridge, MA: Harvard U.P.
- Harman, Gilbert. 1990. The Intrinsic Quality of Experience. *Philosophical Perspectives* 4:31–52.
- Johnston, Mark. 2004. The Obscure Object of Hallucination. *Philosophical Studies* 120:113–183.
- Johnston, Mark. 2006. Better than Mere Knowledge? The Function of Sensory Awareness. In *Perceptual Experience*, edited by Tamar Szabo Gendler and John Hawthorne, 260–290. New York: Oxford University Press.
- Lycan, William. 1996. *Consciousness and Experience*. Cambridge, MA: MIT Press.
- Martin, M. G. F. 2002. The Transparency of Experience. *Mind & Language* 17(4):376–425.
- Moore, G. E. 1903. The Refutation of Idealism. *Mind* 12(48):433–453.
- Nelson, Michael. 2002. Puzzling Pairs. *Philosophical Studies* 108:109–119.
- Noë, Alva. 2003. Causation and Perception: The Puzzle Unravelled. *Analysis* 63(2).
- Peacocke, Christopher. 1992. *A Study of Concepts*. Cambridge, MA: MIT Press.
- Putnam, Hilary. 1973. Meaning and Reference. *Journal of Philosophy* 70(19):699–711.



- Siewert, Charles. 2004. Is Experience Transparent? *Philosophical Studies* 117:15–41.
- Speaks, Jeff. 2007. Attention and intentionalism .
- Stalnaker, Robert. 1984. *Inquiry*. Cambridge, MA: MIT Press.
- Stalnaker, Robert. 1999. Comparing Qualia Across Persons. *Philosophical Topics* 26:385–405.
- Travis, Charles. 2004. The Silence of the Senses. *Mind* 113:57–94.
- Tye, Michael. 1995. *Ten Problems of Consciousness: A Representational Theory of the Phenomenal Mind*. Cambridge, MA: MIT Press.
- Tye, Michael. 2002. Representationalism and the Transparency of Experience. *Noûs* 36:137–151.