How Can We Discover the Contents of Experience?

Citation

Siegel, Susanna. 2007. How can we discover the contents of experience? Southern Journal of Philosophy 45, Suppl.: 127-142.

Published Version

<http://philosophy.memphis.edu/sjp/sjp.htm>

Permanent link

<http://nrs.harvard.edu/urn-3:HUL.InstRepos:3164345>

Terms of Use

This article was downloaded from Harvard University's DASH repository, and is made available under the terms and conditions applicable to Other Posted Material, as set forth at <http://nrs.harvard.edu/urn-3:HUL.InstRepos:dash.current.terms-of-use#LAA>

Share Your Story

The Harvard community has made this article openly available.
Please share how this access benefits you. [Submit a story](#).

[Accessibility](#)

How Can We Discover the Contents Of Experience?*

Susanna Siegel * published in *Southern Journal of Philosophy*, vol XLV, 2006

Suppose you are in a crowded cafeteria, looking at the clusters of people eating and talking. There are elements of the scene that you see - the people, the food, the chairs, tables, trays, and so on – and when you see them, they look some way to you. The *contents* of your visual experience are distinct from the scene in the world that you see. Unlike the things that you see, the contents of your experience characterize how those things look. We can think of these contents as abstract entities – namely, propositions - and of the experiences whose contents they are as being accurate if the contents are true, and inaccurate if the contents are false. For instance, if the table looks longer than it really is, then the experience of seeing it will be inaccurate. And if every way that the scene looks is a way that it is, then the experience will be accurate.

There are many questions we could ask about the contents of visual experience. Most fundamentally, we could ask whether facts about how things look really can be characterized using propositions. If they can, we could go on to ask what kind of propositions are best suited to play this role: Will propositions construed as sets of possible worlds suffice? Or would it be better if they were structured, in the way that Russellian proposition are, by objects and properties? Or can the facts about how things look only be characterized by some sort of Fregean proposition involving modes of presentations of the objects and properties that are seen?

Independently of how we answer these questions, we can also ask what kinds of properties things can look to have when we see them. No matter which kinds of abstract objects turn out to be best for characterizing the way things look, properties will figure in these contents in one way or another: either as properties had by things in the possible worlds that constitute the Russellian proposition; or as constituents of the proposition itself; or as referents of modes of presentation that constitute the Fregean proposition. I will say that a property is represented in experience if it figures in experience contents in any of these ways. And even if the fundamental idea that experiences have propositional contents is wrong, we can still ask what properties things can look to have when we see them. This question have been pursued in debates about whether causation can be perceived (can things look to stand in causal relations, or are putative cases of this really just cases where things look to stand in certain spatio-temporal relations?) and whether

* For helpful discussion thanks are due to Tim Bayne, Alex Byrne, David Chalmers, Tim Crane, Sean Kelly, Fiona Macpherson, John Morrison, Bernard Nickel, Matthew Nudds, Adam Pautz, Rebecca Saxe, Charles Siewert, Maja Spener, Amie Thomasson, Joseph Tolliver, and Michael Tye. A version of this paper was presented to the conference in Glasgow on the Admissable Contents of Experience. Thanks to the conference participants and to the audience at the Spindel conference for their responses.

things can look to be mind-independent.^{1[1]} On one view, the only such properties are very simple. Only colors, shapes and motions of objects and features of their lighting are present to the mind via seeing itself, and the rest of what we come to know by seeing is the result of adding further ingredients to experience, such as beliefs about what one experiences – as it might be, a belief to the effect that a cat is sitting contentedly on a chair. On an opposing view, the properties that things can look to have include more complex ones. In the same way that things can look to have color, shape and illumination properties, they can also look to have properties such as being a cat, sinking into a hammock (note that sinking is a causal notion), or even existing independently of being perceived.

My main concern in this paper is methodological. Rather than trying to answer the question of what properties can figure in visual experience, I will discuss how one might reasonably proceed in answering it. Several natural starting points suggest themselves. First, one might reasonably ask what role, if any, introspection can play in answering the question. Second, over recent years, philosophers including Fred Dretske, Jerry Fodor, Ruth Millikan, David Papineau and others have developed a variety of naturalistic theories that aim to account for what makes it the case that we have contentful mental states, and (in the case of some of these theories) what makes it the case that these states have the contents they do. Some of these philosophers, and others including Robert Stalnaker and Michael Tye, explicitly apply such theories to the special case of experiences. It would be natural to ask whether these applications shed any light on the question of which contents experiences have – on the assumption that they have contents at all.

I will argue that neither the naturalistic theories mentioned nor introspection alone enable us to discover the contents of visual experiences. A different method is needed. I think such a method is available: the method of phenomenal contrast. This paper is an attempt to make explicit what the method is, and to defend it. It is thus addressed to skeptics who doubt that there is any tractable way to find out which properties things can look to have, as well as to enthusiasts who would like to know which properties these are. I am sympathetic to the view that experiences have contents in the sense outlined at the start, and in most of my discussion I'll assume that this view is correct without defending

^{1[1]} Cf. Hume: “All events seem entirely loose and separate. One event follows another; but we can never observe any tie between them. They seem conjoined, but never connected” (1993/1777, Section 7, Part 1) and “as to the independency of our perceptions on ourselves, this can never be an object of the senses; but any opinion we can form concerning it, must be derived from experience and observation (1978/1777, Book 1, Section 2, 191).” Hume could be taken here to be expressing the ideas that pairs of events cannot look to be causally related, and that ordinary objects cannot look as if they exist independently of our perception of them.

it.^{2[2]} Ultimately, however, the case for using the method of phenomenal contrast will not depend on this assumption. The method can be used even if one rejects this assumption, although in that case the initial question the method is designed to answer (“how can we discover the contents of experience?”) would have to be reformulated slightly, in ways that I will mention at the end.

In section 1, I argue that introspection does not help us find out which contents experiences have, and that naturalistic theories of intentionality do not help us find this out in section 2. In section 3, I describe the method of phenomenal contrast and discuss how it may be applied.

1. Introspection

Suppose you are looking at a bowl of expertly designed wax fruits. You have a visual experience when you see this scene, and we can ask which properties your experience represents as instantiated in the scene before you. Here are two hypotheses. According to the color-shape hypothesis, your visual experience represents the colors and shapes of the wax fruits, but does not go so far as representing that they are fruits. Moreover, the wax fruits actually have the colors and shapes that the experience represents them as having. If the color-shape hypothesis about the contents of experience is correct, then the experience is veridical, because it does not attribute to the wax fruits any properties that they lack. You may go on to form a false belief to the effect that there are cherries in the bowl, but if you do that, your error will remain at the level of belief, and will not infect your experience, since your experience is not committal on whether the colored shapes in the bowl are cherries or not.

According to the cherry-content hypothesis, your visual experience represents the property of being a cherry, and its contents include that there are cherries in the bowl. If the cherry-content hypothesis is correct, then your experience is falsidical, and in having it you have a visual illusion. In contrast, if the color-shape hypothesis is correct, then even if you end up with a false belief that there are cherries in the bowl, you won’t have a visual illusion, because there will be no error at the level of your visual experience.

Could you decide between these hypotheses, just by introspectively reflecting on your experience of seeing the bowl of wax cherries? It seems plain that in seeing the fruit bowl (or even merely in hallucinating), one can know by introspection that one is not having an experience as of a busy airport, or an undifferentiated expanse of blue. Introspection can rule out many proposed contents as inadequate to the phenomenal character of the experience, and thus does not seem to be completely useless as a means of finding out which contents an experience has. It thus seems to take us some way toward finding out which contents a specific experience has. But it does not take us far enough. Both the color-shape hypothesis and the cherry-content hypothesis seem *prima-facie* plausible – neither is obviously phenomenally inadequate. If introspection could tell us which of them were correct, then it would tell us directly about the phenomenal

^{2[2]} I defend the view in *The Contents of Visual Experience*, forthcoming from Oxford University Press.

character and content of experience to a degree of precision that would be needed to decide between these hypotheses. And if it did that, one of the hypotheses *would* seem obviously phenomenally inadequate. Since it doesn't, we can conclude that introspection does not decide between these hypotheses.

It may seem strange that introspection could not decide between these two hypotheses, since the contents of experience are supposed to reflect the phenomenal character of the experience. Moreover, states with phenomenal character may seem to be intrinsically the sort of states that are accessible to introspection.^{3[3]} But even if we grant that phenomenal character is intrinsically accessible to introspection, it does not follow that introspection tells us sufficiently precisely which contents such states have. For one thing, though it is sometimes straightforward enough to elicit reports that are based on introspection – one might report that there look to be cherries in the bowl, for instance – for these reports to tell us what contents experiences have, we would need to have principles linking reports to experiences, and such principles are elusive. The simple principle that identifies the contents of the reports with the contents of experiences faces the difficulty that there is no natural language expression used exclusively for reporting the contents of experiences as opposed to the contents of beliefs that one actually has formed or is disposed to form on the basis of experience. So some more complex principle would be needed to link such reports and contents. The fact that a report is based on introspection of an experience does not suffice to show that its contents are experience contents.

Another difficulty with relying exclusively on introspection is that it is not clear exactly what procedure a person would follow, if she wanted to use introspection to find out which contents her experience had. This difficulty comes in a practical and a theoretical version. The practical version is relatively innocuous: it is unclear what would count as following an instruction to introspectively reflect on one's visual experience. Although this unclarity may prevent an individual from easily assessing whether introspection alone can find out what contents experiences have, it doesn't by itself refute the proposal that introspection is indeed capable of finding this out. Merely having a practical difficulty discerning what kind of reflection is introspective reflection is compatible with there being such a thing as distinctively introspective reflection, and with that kind of reflection informing us of what contents experiences have.

The theoretical version of the difficulty, in contrast, poses a problem for relying exclusively on introspection to decide between competing hypotheses about the contents of experience – or to decide from scratch, without initial hypotheses, which contents these are. The theoretical difficulty concerns what kind of procedure or faculty introspection is. On a domain-general model, introspection is a way to find out about the current occupants of one's stream of consciousness, whatever they are; but it may not be

^{3[3]} This issue is at the center of the controversy over whether there can be phenomenal consciousness without what Ned Block calls "access-consciousness". See Block, "Consciousness, Accessibility, and the Mesh Between Philosophy and Neuroscience", forthcoming in *Brain and Behavioral Science*.

introspectively obvious whether an occupant is a visual experience, as opposed to a different kind of mental state or event.^{4[4]} If this model best reflects the nature of introspection, then the strategy of relying exclusively on introspection to decide between hypotheses about the contents of experience faces the pitfall that one might unknowingly introspectively attend to something other than visual experience, even while introspecting correctly. If so, then verdicts delivered by introspection that purport to be verdicts on the contents of visual experience may in fact be about the contents of some other kind of state, such as a judgment.

If the domain-general model of introspection is not apt, this problem could not arise. An alternative model of introspection is a domain-specific model. On a domain-specific model, introspection of visual experience is a different kind of reflection (perhaps because it is a different kind of procedure) from introspection of other kinds of occupants of the stream of consciousness, so there is no possibility of mistaking a visual experience for some other kind of occupant of the stream of consciousness, or of making the opposite mistake.^{5[5]}

It is doubtful that the domain-specific model is apt, if introspection is construed as a procedure whose employment is an optional addition to the experience itself. There doesn't seem to be any special sort of procedure of reflection that cannot be used for visual experience but can be used on, say, olfactory experience. On the assumption that the procedure is a kind of attention, for example, there seems to be only one kind of attention at the person-level, which can be variously directed to what is seen or what is smelled. A piece of evidence for this is that directing attention in one of these ways can interfere with directing attention in the other.

The only way in which introspection could plausibly be domain specific is if it were construed as a perspective on experience that was somehow built in to visual experience. Either such a perspective would often be inaccurate in its portrayal of the contents of experience (even when there were no intrinsic defects in the experience) or it would not. If it were inaccurate, then it could not be used to settle which contents experiences have. If such a perspective were reliably accurate, then in principle it would indeed be able to settle such disputes. But we do not in fact have any direct line of access from this perspective to what we report when we reflect on our experiences, since we find ourselves initially uncertain about exactly which properties experiences represent. Either way, then, introspection on its own, construed as a perspective on how things look that is

^{4[4]} The domain-general model need not rely on the idea (which some find suspicious) that there is such a thing as a stream of consciousness with occupants. It could equally be considered a procedure that purports to discern that one is currently in a mental state, without discerning exactly what kind of state it is.

^{5[5]} Like the domain-general model, the domain-specific model also need not rely on the idea that there is a stream of consciousness with episodic occupants. Instead, it could be construed as a procedure for attending to features of the state of visual experiences that one is in when one is in it.

built into our experiences, does not enable us to discover the contents of those experiences.

In the next section, I consider whether some influential naturalistic theories of content for mental states might be employed to decide which contents experiences have.

2. Naturalistic theories of content

According to some philosophers, for a mental state to have content is for it to represent that such-and-such conditions obtain, where representation is a kind of causal co-variation with those conditions. Such theories are controversial. I will argue that even if they are true, they will not help us discover the contents of experience.

Michael Tye has developed a causal-covariation theory for visual experiences.^{6[6]} He states the theory as follows:

For each state S of object x, within the relevant set of alternative states of x:
S represents that P = If optimal conditions obtain, S would be tokened in x if and only if and because P.^{7[7]}

Robert Stalnaker endorses something close to this theory as well.^{8[8]}

One might think that the causal co-variation theory would make straightforward predictions about which contents experiences have, by licensing the following procedure: pick an experience, identify the optimal conditions for having it, and find out what causes and co-varies with the experience under those conditions. Call this the *experience-first procedure*. If the causal co-variation theory is true, then – one might think – the content of the experience would be the proposition that such-and-such obtains, where such-and-such is what causes and co-varies with the state. We would then either have to take these predictions as the truth about what the contents of experiences are, or else refute the theory.

To follow the experience-first procedure, one would have to identify which experiences in counterfactual circumstances (namely, other optimal conditions) are the same as the one we were asking about. This is not something the causal co-variation

^{6[6]} Tye calls these ‘simple perceptual sensations’ which are supposed to be outputs of what he calls ‘specialized sensory modules’ (op cit. p. 103).

^{7[7]} Tye op cit p. 101. As we will see, this is only part of the theory. A separate part tells us what makes such a state an experience.

^{8[8]} “What Might Non-Conceptual Content Be?” in Gunther, ed., *Essays on Non-Conceptual Content*. MIT Press.

theory itself tells us.^{9[9]} We would need a way to identify which experiences across possible worlds are the same, holding constant the essential features of the experiences, whatever they are, while allowing others to vary. Without a way to do this, we would not be able to assess whether the experience whose contents we are asking about *would be tokened* in x. If we can't assess that, then we can't consider what situations would cause and co-vary with such counterfactual tokenings.

Before we have settled what contents the experience has, however, it is difficult to reidentify experiences across worlds. It might be plausible to suppose that if my current experience has content that p, then any experience in some other possible world that's identical to it will also have that content – but which contents my experience has is just what we're trying to find out. Moreover, our initial uncertainty about which contents experiences have is also uncertainty about their exact phenomenal character. The phenomenal character is not more accessible to introspection than the contents of experience. We are thus also precluded from using the phenomenal character to reidentify experiences.

There is a different potential procedure for discovery, however, that does not face this difficulty. Instead of beginning with an experience and trying to identify it across possible worlds, the alternative procedure begins with some other state – as it might be, a neural state, or a functional state of some sort – initially ignoring the question whether this state has the status of an experience (or indeed of any mental representation). Call this the *information-state-first procedure*. The next step is to try and find out what conditions are optimal for being in this state, and what causes and co-varies with the state under those conditions. Finally, with the informational content of the state in hand, one consults the theory's criterion for being an experience, as opposed to being merely subpersonal state or a mental state of some other sort. For Tye, such a state is an experience if at least some of its contents have the proper functional role in fixing belief, and Stalnaker makes a similar albeit more tentative suggestion.^{10[10]} (The information-state-first procedure thus mirrors the structure of the causal co-variation theory). The contents of the informational state that are poised would then be the contents of an experience – namely, the experience that the informational state realizes.

The limits of this method become apparent when we consider what relation the neural state might stand in to the experience. Either it underlies the experience in the

^{9[9]} It need not be a fault of the causal co-variation theory that it is silent on how any mental states it applies to are identified across worlds. The causal co-variation theory was developed in an effort to show that mental states could represent happenings in the world, consistently with physicalism. If the theory is correct, then it seems to succeed in doing that, even if it does not tell us how to identify experiences across worlds. Even if for some reason the theory turned out not to be consistent with physicalism, it could still be correct in holding that mental states represent what they causally co-vary with under optimal conditions.

^{10[10]} Stalnaker op cit, p. 103, Tye op cit chapter 5.

manner of a realizer (let's consider being identical to it to be a special case of this), or it is only part of the realizer of the experience. Only in the first case will we have learned what contents the experience has from learning what contents the neural state has. For the method to be useable and informative in such cases, our epistemic position at the start has to be this: we have good reason to think that the neural state is the complete realizer of the experience, but we don't know what contents the experience has. Since we don't now have much understanding of what sorts of neural states exhaustively realize experiences, we can't actually use this method - even if the causal co-variation theory is true.

Now suppose the informational state (neural state, or a functional state) is only a partial realizer of the experience. Then if we apply this method, we will have learned (by the lights of the causal co-variation theory) something about what contents the experience includes. This would be a partial discovery about the contents of experience. We won't, however, have put any limits on which contents these can be, and so applying the method in this sort of case may not always help decide between the hypotheses of the sort discussed at the start.

We are often in a similar epistemic situation when we learn about mechanisms or processes that underlie certain kinds of conscious experiences. Suppose it turned out that the mechanism or process underlying gestalt groupings was the same as the one used to track objects.^{11[11]} It may be reasonable to conclude that our experience is veridical only if there is a single object that has the features that the grouping looks to us to have. But as in the case where we know the informational content of a neural state that only partially realizes the experience, we can't conclude that the contents of experience are exhausted by the kind of information carried by the mechanism or process underlying the gestalt grouping. For all this result says, experience may also be able to represent that the 'object' has causal properties, or that it belongs to (or depicts something that belongs to) a certain kind.

Finally, even if some application of the method derived from the causal co-variation theory can tell us about the contents of an experience, it does not let us start with the experiences we want to know about, such as the experience of seeing faces, or experiences that controversies disagree about.^{12[12]} Instead, we have to happen upon any

^{11[11]} Thanks to Matthew Nudds for the example. If the mechanism underlying gestalt groupings is an attentional mechanism, then it probably isn't the same as the mechanism underlying tracking of objects, since the latter has limits on the number of objects that can be tracked, and such limits do not seem to restrict what we can attend to.

^{12[12]} The experiences of seeing faces is important in understanding the nature of certain pathologies of belief, such as Capgras syndrome, in which subjects believe that their spouse has been replaced by an impostor. Discussion of the nature of this syndrome focuses on whether subjects are responding (normally or abnormally) to a normal experience, or to an abnormal one. To settle this, we need to know what experiences of faces normally represent.

experiences whose contents we want to know about indirectly, by starting with an informational state.

The method derived from the causal co-variation theory thus has its limits. I will now argue that another method is available that can produce better results.

3. The method of phenomenal contrast

The method of phenomenal contrast is a way to test a target hypothesis about the contents of experience. The main idea behind the method is to find something that the target hypothesis purports to explain, and see whether it provides the best explanation of that phenomenon. Instead of taking a specific experience as input and delivering as output a verdict on its contents, then, the method's starting point is a target hypothesis, and it aims to reach a yes-or-no verdict. It is thus a way of testing hypotheses, rather than a way of generating hypotheses in the first place.

Since contents are supposed to be phenomenally adequate, any target hypothesis will predict that any two experiences differing with respect to the hypothesized contents will differ phenomenally as well. It is thus possible to use the phenomenal contrast as the thing to be explained. The 'target explanation' will say the experiences contrast phenomenally because one of them has the hypothesized contents, while the other one does not.^{13[13]}

Several live debates about perception can be construed as employing the method of phenomenal contrast. Consider the supervenience thesis (sometimes labeled 'intentionalism') that any two experiences with the same representational properties have the same phenomenal character. Some philosophers argue that cases of size and color constancy show intentionalism to be false. In response, intentionalists try to find differences in representational properties accompanying the phenomenal contrast. In his 1983 book *Sense and Content*, Peacocke discussed a case of seeing two trees of the same size on a straight street. Peacocke claimed that we don't suffer any illusion about whether the trees are the same size (they look to be the same size), and yet, he claimed, there is a sense in which the tree that is farther away looks smaller than the tree that is closer. Much debate ensued about how best to account for the phenomenal contrast between the two trees – whether it was a difference in content, and if so, how the content of the experience of seeing each tree differed. Here is an example of a phenomenal contrast serving as an explanandum which hypotheses compete to explain.

The same argumentative role is played by the phenomenal difference between experiences of spectrally inverted perceivers. Such perceivers allegedly have phenomenally different visual experiences (one is phenomenally like a normal experience of seeing unpeeled cucumbers and the other is phenomenally like a normal experience of

^{13[13]} If the target hypothesis does not figure in the best explanation of the phenomenal contrast, then the method (properly applied) will reveal this. But it will not necessarily identify which explanation of the phenomenal contrast is the best. So although the method is a version of inference to the best explanation, it is not a version that always yields a verdict on which explanation is the best.

seeing ripe strawberries), yet stand in the same kind of causal relation to the physical property identified (for the sake of the thought-experiment) with redness.^{14[14]} Subsequent debate has revolved around whether the phenomenal difference is a difference in content, and if so, what contents each experiences have.

Both of these debates involve phenomenal contrasts that various positions compete in trying to explain. The phenomenal contrasts involved in perceptual constancy and in putative spectral inversion pose long-standing challenges to psychology and philosophy, and as such are already salient.^{15[15]} But the method of phenomenal contrast can also be used to test hypotheses that single out a specific property, such as the cherry-content hypothesis, or the hypothesis that causal relations can be represented in perception. These hypotheses are more specific than intentionalism. When the method is employed to test them, a pair of contrasting experiences has to be sought out. The contrasts that figure in the applications of this method are not likely already to be subjects of debate. Finding the right contrast is itself difficult for a proponent of a target hypothesis, since from the perspective of such a proponent the contrast has to survive as an illustration of the contrast between an experience that illustrates the hypothesis and an experience that does not.

We can illustrate how the method works by considering the hypothesis that visual experiences can represent the property of being caused by something. Let this be the target hypothesis in our example. Target contents will be contents involving causal properties. We can then consider a pair of experiences with the following two features.

- ♦ First, one of the experiences in the pair, call it the target experience, is a reasonable candidate for an experience that represents causation (the target property). An example for the case of causation might be an experience in which one flicks a light switch and sees the light go on. This experience is a candidate for one that represents that the flicking of the switch caused the light to go on. In contrast, the other experience should clearly and uncontroversially *not* represent the target property (causation). Here one might consider a similar sequence in which the light that turns on is at a great distance from the switch that is flicked.
- ♦ Second, the experiences in the pair fairly obviously differ phenomenally, so that people who may ultimately disagree about the target hypothesis would nonetheless agree that they exhibit a phenomenal contrast.

^{14[14]} For discussion of intentionalism and the inverted spectrum, see the Introduction in Byrne and Hilbert (1997) and the papers therein.

^{15[15]} On spectral inversion, see J. Locke, *Essay Concerning Human Understanding*. 1689/1975, *Essay Concerning Human Understanding*, Oxford: Oxford University Press, and Nida-Rumelin. On constancy, see Russell (1997/1912), Smith (2002).

If the target hypothesis figures in the best explanation of the phenomenal contrast, it will be true of the target experience but not of the contrasting experience. According to this explanation, the reason there is a phenomenal contrast between the experiences is that one of the represents the target property while the other one does not.

To use the method to decide whether the target hypothesis figures in the best explanation of the phenomenal contrast, the target explanation has to be compared to alternative explanations. Which alternatives should be considered? We saw one type of alternative earlier when we considered the fruit-bowl example: there was the cherry-content hypothesis on the one hand, and the color-shape hypothesis, on the other. More generally, whereas the target explanations I've described so far posit a difference in content, competing explanations either posit a different difference in content, or they posit a difference of some other sort.

We can divide up the types of alternatives that matter by focusing on three general commitments that are entailed by the claim that the phenomenally contrasting experience differ with respect to whether they represent F. By considering challenges to these different commitments, we can arrive at three types of competing explanations of the phenomenal contrast.

The following three claims will be true if the target experience represents F while the contrasting experience does not:

- ♦ First, the target experience will differ in its sensory phenomenal character from the contrasting experience, rather than differing merely with respect to a non-sensory category of phenomenal character.
- ♦ Second, the target experience will have representational content.
- ♦ Third, the target experience will have a representational feature – namely, representing the property F - that the contrasting experience lacks.

These three general commitments indicate three types of alternatives that challenge the target explanation.

- **Alternative 1: No difference in sensory phenomenology**
The phenomenal contrast is exclusively a difference in non-sensory phenomenal character.

Since the target explanations purport to explain differences in visual phenomenology, if the only phenomenal difference is a difference in non-sensory (a fortiori non-visual) phenomenology, then they lose their explanandum. In assessing alternatives of type 1, one may proceed in general terms, by defending a way to distinguish sensory phenomenology from other kinds. Alternatively, one may proceed without making such general commitments, focusing on the specific experience at issue, and the plausibility in that case of the view that the relevant bit phenomenal character is sensory as opposed to non-sensory.

- **Alternative 2: No difference in content**

The two experiences do not differ at all in their representational content, a fortiori, they do not differ with respect to representing F.

This type of challenge to target hypotheses encompasses two different sorts of alternatives: the experiences in the pair could fail to differ in content either because neither has any content in the first place, or because the content each has is the same. Either way, something other than a difference in visual experience content is invoked to explain the phenomenal contrast. Perhaps the simplest example is a raw feel or 'quale'. Proponents of qualia who appeal to this type of alternative have two importantly different theories of the structure of visual experience available to them. First, according to a pure qualia view, there is nothing to visual experience other than raw feels or qualia. Second, according to a two-factor view, visual experience consists of a combination of qualia and dispositions to form beliefs about what one seems to see, such as beliefs that one thing causes something else to happen.

Since even a proponent of the pure qualia view will presumably grant that subjects of experiences have dispositions to form such beliefs in addition to qualia, the difference between these two views may seem merely terminological, as it concerns which cluster of ingredients experiences are identified with, rather than which ingredients there are in the first place. However, the difference between these views is not merely terminological. A pure qualia view has only one ingredient to invoke to explain phenomenal contrasts. A two-factor view has two: they may appeal either to a difference in qualia, or a difference in dispositions to believe that by their lights are part of the experience. For instance, in the causation example, a two-factor view might hold that the qualia are the same in both experience, but the phenomenal contrast is due to a change in whether the subject is disposed to form causal beliefs.

It may at first seem odd to consider a difference in what the subject is disposed to believe to be a non-representational difference. But what's meant in alternative 2 by 'non-representational difference' is a difference in the contents of experience, where contentful experiences are construed as propositional attitudes involving an attitude of experientially entertaining a content that is distinctive of visual experience. So it is built in to the target theses as we've been discussing them so far that they are incompatible with two-factor views according to which there is no such distinctive attitude.

There is nothing in the method of phenomenal contrast, however, that rules out its use as a means to test hypotheses that are more liberal about what shall count as contents than the target hypotheses so far introduced. If one allowed experiences structured by the two factors (qualia and dispositions to believe) to count as contentful experiences, then target hypotheses would face fewer challenges. But since the differences are substantial between two-factor views and the target hypotheses initially introduced, the method potentially tells us more about the structure of visual experience if it preserves their differences.

Since alternatives of type 2 concern the structure of visual experiences, in assessing such alternatives one may proceed in general terms, by considering the

plausibility of the general view about the structure of experience, such as the pure qualia view, or the two-factor view. However, it is also possible not to proceed in entirely general terms. Since the method involves considering a specific phenomenal contrast, the key question is whether the specific aspect of phenomenal character that it present in the target experience but absent in the contrasting experience is plausibly nothing but a raw feel.

- **Alternative 3: Difference in content**

The two experiences differ in their representational content, but neither represents the target property.

The difference between the cherry-content hypothesis and the color-shape hypothesis is an example of this type of alternative. Likewise, an alternative of this type in the causation case would say that instead of differing in whether they represent causal properties, the two experiences in our example differ in whether they represent certain color and shape properties.

On what basis could one decide between target explanations and the three types of alternatives? If we could rely only on introspection to decide between them, then the method of phenomenal contrast would face the same difficulties as the method of introspecting single cases. Hardly anything general can be said about the basis on which to decide between the alternatives at this stage, since the considerations will vary from case to case. In some cases one might be able to consider specific features of alternative proposed contents, such as whether they posit either indeterminacy or mental entities to account for the phenomenal contrast, and see whether in those specific cases it is plausible to do so.^{16[16]} Or one might consider analogous phenomenal contrasts and see whether it is plausible to suppose that the proposed explanations would apply in those cases as well.^{17[17]}

The method of phenomenal contrast is a way to limit the use of introspection in theorizing about visual experience. All that introspection is relied upon to do is to detect the phenomenal contrast. The method need not take a stand on the *category* of the phenomenally contrasting states, such as whether they are sensory, cognitive or some other kind. It also need not take a stand on the *structure* of the phenomenally contrasting states, such as whether either is fundamentally a state with content, as opposed to a contentless state, or to a two-part state. Finally, introspection need not take a stand on

^{16[16]} The example of indeterminacy is taken from the fifth alternative to the target hypothesis canvassed in Siegel (2006) p. 377, and the example of mental entities from the third and fourth alternatives, op cit p. 376.

^{17[17]} One example of this sort of consideration is discussed in Siegel (2005), in connection with a case of seeing a contorted face before and after you learn that it is expressing doubt. Another is discussed in Siegel (forthcoming), in connection with the Nielson experiments.

whether the phenomenal difference between the experiences in the pair is also a difference in content, let alone on what the difference in content, if there is one, may be.

In describing the method of phenomenal contrast, I have focused on an application of the method to test a target hypothesis according to which the contents of experience are ‘rich’ as opposed to impoverished. But nothing in the method itself biases it toward this conclusion. I have also focused on just one type of application of the method - the application to target hypotheses concerning contents of experience, when contents are construed as accuracy conditions. The method itself, however, is simply a structure for testing hypotheses that doesn’t put any constraints on what the hypotheses say – let alone which hypotheses are true. The same method could be used to answer the question of what contents experiences have, when contents are construed in some other way. For instance, consider a possible target hypothesis according to which you can be ‘appeared-to’ causally.^{18[18]} Given the theoretical role of properties of being appeared-to, there will be a pair of experiences that contrast phenomenally, where by the lights of the target hypothesis, one of the experiences is a being-appeared-to-causally, and the other isn’t. The method of phenomenal contrast could be used to test this hypothesis against alternatives. One could use the same structure as the one described to identify the kinds of alternative hypotheses that would have to be considered. In place of ‘no difference in content’ at stage 2, one could substitute ‘no difference in ways of being appeared-to’, and in place of ‘a different difference in content’ at stage 3, one could substitute ‘a different difference in ways of being appeared-to.’ Analogous adjustments could be made for views that deny that experiences that deny that experiences have contents construed as accuracy conditions.

Works Cited.

Block, N. (forthcoming) “Consciousness, Accessibility, and the Mesh Between Philosophy and Neuroscience”, in *Brain and Behavioral Sciences*.

Byrne, A. and Hilbert, D. (1997) *Readings on Color*, vol. 1. Cambridge: MIT Press.

Dainton, B. (2000) *Stream of Consciousness*. London: Routledge and Kegan Paul.

Ducasse, C. J., 1942. “Moore's ‘The Refutation of Idealism’” in P. Schlipp (ed.), *The Philosophy of G.E. Moore*. Evanston, IL: Northwestern Press.

Hume, D. (1993/1777). *An Enquiry Concerning Human Understanding*. Second Edition. Indiana: Hackett Publishing Company.

^{18[18]} The notion of being appeared-to in some way is employed by Ducasse (1942), a champion of adverbialism about visual experience.

----- (1978/1777) *A Treatise of Human Nature*, ed. P. H. Nidditch. 2nd ed.
Oxford: Clarendon.

Johnston, M. (2004). "The Obscure Object of Hallucination" *Philosophical Studies* 103:113-83.

Nida-Rumelin (1996) "Pseudonormal Vision: An Actual Case of Qualia Inversion?" *Philosophical Studies* 82, 145-57.

Peacocke, C. (1983) *Sense and Content*. Oxford: Oxford University Press.

Russell, B. (1997/1912) *The Problems of Philosophy*. New York: Oxford University Press.

Siegel, S. (2005) "Which Properties are Represented in Perception?" in T. Gendler and J. Hawthorne, Eds, *Perceptual Experience*. Oxford: Oxford University Press.

----- (forthcoming) "The Phenomenology of Efficacy", *Philosophical Topics*

----- (2006) "Subject and Object in the Contents of Visual Experience" *Philosophical Review* 115:3.

Smith, A. D. (2002) *The Problem of Perception*. Cambridge: Harvard University Press.

Stalnaker, R. "What Might Non-conceptual Content Be?" in Gunther, ed. *Essays in Non-conceptual Content*. Cambridge: MIT Press.

Tye, M. (1991) *Ten Problems of Consciousness*. Cambridge: MIT Press.