

Phenomenal Properties: Some Models from Psychology and Philosophy

Austen Clark, Department of Philosophy, University of Connecticut
13 May 2008. Forthcoming in *Philosophical Issues*, vol 18.

Abstract. There are many different notions of "phenomenal property", and these notions differ markedly in logical, semantic, and metaphysical complexity. This paper identifies and analyzes a traditional and relatively simple notion of phenomenal properties, characterizing how the world appears to the senses. I argue that developments in philosophy and experimental psychology have made this notion of phenomenal property eminently precise and empirically tractable. Psychologists have had great success in constructing explanatory models of various modalities of sensory appearance. The strategy and structure of such psychological model building is analyzed in some detail. One element critical to the success of that enterprise is that psychologists have (in effect) confined themselves to explananda describable in the comparative (or relational) forms of the verbs of appearance, while philosophers continue mostly to obsess about the non-comparative senses. The paper shows how five features of philosophical discourse about phenomenal properties can be fruitfully redescribed in relational terms. In the concluding section this traditional and tractable notion of phenomenal properties is contrasted with a more complicated and much less tractable kind.

Acknowledgments. I thank Alex Byrne, Mohan Matthen, and Murat Aydede for comments on earlier versions of this paper, as well as audiences at Harvard University, the University of Toronto, the University of British Columbia in Vancouver, and Rutgers University.

In most domains philosophers struggle mightily with Reality. Philosophy of perception is one of those exceptional domains in which even murkier and more difficult struggles lie: those that grapple with mere Appearance. This wrestling match began in ancient times and results so far are inconclusive. It has proven difficult to pin an opponent who merely seems to be present.

In this paper I want to describe one line of development, occurring mostly within the twentieth century, that identifies one subset of opponents within this competition against whom some progress is *possible*. It is a rather important and central subset, but there is no denying that it is a subset; other denizens of appearance are left untouched by the stratagems herein proposed. Nevertheless I think it *is* progress, worth noting and analyzing, to separate apparent wheat from apparent chaff.

In particular I will analyze one relatively simple kind of "phenomenal property", distinguish it from some of its more complex cousins, and show how this simpler kind has proven to be a tractable target for empirical investigation in experimental psychology. The cost of opening a notion to empirical investigation is that elements of the conventional wisdom about that notion might be empirically disconfirmed, and I shall argue that this is precisely what has happened with this tractable variety of phenomenal property. In particular, I shall argue, these properties have been and must be firmly dissociated from any ties to consciousness.

I

Etymologically, "phenomenal" has a root meaning of "that which appears or is perceived by the senses" but it has acquired other technical connotations in philosophy of mind through its associations with various doctrines of phenomenalism, phenomenology, and phenomenal consciousness. I shall initially set these latter accretions aside, and focus on the simplest, root notion. The latter can trace its lineage back to ancient worries about the distinction between appearance and reality; specifically the distinction in perception between how things seem and how they are. A "phenomenal property" is, very roughly, a characteristic of sensible appearance: a quality that qualifies how things *appear*. A classic source of examples is the study of perceptual illusion: of things seeming to have perceptible qualities that they do not have. It bears emphasis that in this root notion it is *things* that appear, and they must appear to the senses. The notion is not meant to characterize things as they are represented in thought, for example; nor is it a higher order notion characterizing one's perceptions or sensations of things. In the first instance phenomenal properties characterize how the *world* seems: how the entities that one perceives in the world *appear* to the subject who perceives them.

The verbs "seem" and "appear" are two out of many of what have been called "verbs of appearance", and it is useful immediately to distinguish some of the different uses of these verbs (see Chisholm 1957). One use that is not particularly germane to the concerns of this paper is the "epistemic" use, in which "The train appears to be moving" (or "feels as if it is moving" or "sounds as though it is moving", etc.) is taken to imply that the speaker (a) believes that the train is moving, and (b) has perceptual evidence of some variety (which variety perhaps being indicated by the particular verb) which the speaker takes to be adequate for that belief. Here "appears" "looks" "seems" and so on are more or less synonymous with "evidently" or "apparently"; the speaker is making an epistemic claim, about the availability of perceptual evidence for a particular conclusion.

The "non-epistemic" or "phenomenal" senses of these verbs do not require or imply that the subject has any beliefs, and indeed the most interesting cases are precisely those in which (for example) the train appears to be moving even though one knows perfectly well that it is stationary. Such an experience might initially tempt one to the mistaken judgement that the train is moving, but continue even after one has realized that such a judgement would be a mistake. So such an experience must be something distinct from a mistaken judgement (see Broad 1927, 236).

Chisholm (1957) identified two non-epistemic uses. In the first, "comparative" sense, sentences of the form " x appears to S to be F " mean that x appears to S in the way in which things that *are* F appear under conditions that are so-and-so. The conditions "so-and-so" must differ from the ones that are current (the ones in which the non- F thing appears to be F), and they are understood to identify some circumstances in which things that are F appear as

they really are, but otherwise they are not specifically identified. Chisholm argued that understanding and (in particular) justifying sentences that use the verbs of appearance comparatively requires a "non-comparative" sense, of the form "*x* appears *F*", whose justification does not require appeal to any incidental information from comparisons across perceptual occasions. In his famous example, a harried witness is pressed to defend his claim that the cabin on Mt. Monadnock looks to him just like things that are blue look under standard conditions. His justification is (a) The cabin looks blue, and (b) things that look blue look just like things that are blue look under standard conditions. The (a) claim must be understood noncomparatively. If it in turn is questioned, the witness first pushes all claims about objects external to his experience to the "incidental information" side (eg "something looks blue, and that something is a cabin") and then finally, to justify the first of these, shifts into a passive-form characterization of what is experienced: "I am being appeared-to blue" or "I am aware of blue" or "I sense blue".

An initial (mid-twentieth century) characterization of phenomenal properties is that they are the truth-makers (if there are truth-makers) for sentences employing the non-epistemic and non-comparative sense of verbs of sensory appearance.

A simple example will help. Some of the early "energy-saving" light bulbs have poor color rendering characteristics. Their fluorescent coatings have emission spectra with several sharp peaks in the green part of spectrum and not much in between. The result is a green cast to the light, much more noticeable with some objects rather than others. Skin tones and complexion often suffer, for example. Another is that some kinds of olive oil look vividly (and disturbingly) green. Employing this energy-saving light may help save the world, but the olive oil in it will look green. Going by the way it looks, you would judge it to *be* green. But you know that it is not. We can mimic John and Jim, protagonists in Sellars' story of the necktie shop (see Sellars 1963, 143). "I don't know what to say", said John, looking at what seems to be--but is known not to be--a green necktie, lying in the newly installed electric light in his shop. Jim taught him to suppress the copula "is", and substitute "looks". The same strategy works today. It *looks* just like olive oil would look if it were green, and it looks just the way things look when I judge them to be green, but something in these circumstances is amiss, and I wish to withhold that judgement.

Now if we think of a phenomenal property as some property borne by some entity such that "the olive oil looks green" is true in virtue of that entity bearing that property, then it is quite clear that the bearer of that property is not the olive oil *per se*.

Suppose we undertake a Chisholm-style interrogation analogous to the one about the cabin on Mt. Monadnock. What reason do you have for saying that the olive oil looks green in this light? We first get to a claim which makes a non-comparative use of the verb of appearance: "It is because something (which is, in fact, the olive oil) appears green [noncomparatively] and I know

that things that appear green noncomparatively in this light look green in this light". And one further step (to "what reason do you have for thinking that something appears green?") drops the "something" from "something appears green" and arrives at the passive form of appears-to: "I am being appeared-to greenly, and I know that if I am appeared-to greenly and ... [e.g., the other incidental information obtains] then something appears green."

This "I am being appeared-to greenly" is exactly the same claim one would make in a situation in which the claim to be defended is that something in the vicinity *is* green. Here though the "independent information" is such that being aware of green does not suffice to justify the claim that something *is* green; it justifies only the claim that something, in this viewing situation, *looks* green. Now obviously the difference between those two situations can't be described solely by reference to the object in question. In one situation some claim of the form "Something appears green" serves to justify the claim that the olive oil *looks* green. In another situation a claim of the form "Something appears yellowy amber" serves to justify the claim that the very same olive oil *is* yellowy-amber. Since it is the same olive oil in both situations, the olive oil *per se* cannot account for the differences between these situations. The difference in appearance must arise from other variables, present in one situation but not in the other.

II

How might we explain these differences (and similarities) in appearance? To state the explanandum we use a comparative sense of a verb of appearance. Why does the olive oil in energy-saving light look the same as something green seen in daylight? Here we use "looks the same as", but it has many variants. For other modalities we will need "feels the same as", "smells the same as", etc. More generalized variants include "appears the same as", "matches", "resembles", "is indiscriminable from", and other relations out of the set of what I will call "relations of phenomenal similarity", including relative forms (e.g. x is more similar to y than to z , x is more similar to y than w is to z , etc); and the contraries or negations of all these ("phenomenal dissimilarities"). Then "looks green", if needed, can (as noted by Chisholm) be represented as "looks (in these circumstances, to that subject) the same as something that (to that subject, in such and such circumstances) is green". The right side ("is green", "is red", "is blue", and so on) is going to require some sort of analysis in any case. In one fell swoop this maneuver eliminates from the philosophical "to do" list the need separately to analyze a mass of distinct phenomenal properties ("looks green", "looks red", "looks blue", and so on). Instead, we have just some relation of phenomenal similarity (here "looks the same as") and the standard color predicates, ascribed in standard ways.

C. D. Broad (1927, 237) urged upon us the need to state the "facts of sensible appearance" in a perspicuous form, begging as few questions as possible, so that theorists of different stripes could at least agree upon what needs to be explained. If we confine ourselves to listing what matches what,

what is similar to what, and what is relatively dissimilar from what, we would be close to this ideal.

Matching gives something similar to what Carnap (1967, 22) called a "pair list", though what matches are not two *elementar Erlebnisse*, and the relation is not "recollection of similarity". Instead the two entities that look the same are probably best thought of as two portions of the world, perceived by some particular subject in some particular circumstances. It is not true that the olive oil per se looks green; instead it is the olive oil in these particular viewing circumstances, seen by a particular observer. If we embrace *scenes* as primitives for purposes of perceptual theory (Barwise 1981; Barwise & Perry 1983) then the x that looks green is, precisely, some portion of a scene. Like a stimulus, numerically the same scene cannot be presented twice. Psychologists might more happily think of it as a portion of the ambient optic array: some subset of the totality of visual stimuli available to a particular subject over some particular interval. That aspect of a thing, portion of the scene, or portion of the ambient optic array is what looks green.

The portion is delimited spatially and temporally. Various aspects of the viewing conditions affect the visual appearance of a target, including the stimuli surrounding it, the state of adaptation of the observer, the viewing angle, the orientation of the target, and a laundry list of conditions of illumination. The latter include not only properties of direct lighting but also the effects of light scatter and reflection from surrounding objects. The list of parameters for "viewing conditions" has as many entries as there are parameters in the perceptible situation that can affect the appearance of the thing perceived. It might require some m parameters to specify, so I will call the terms naming these entities " m tuples".

If we include lists of terms satisfying all the relations of phenomenal similarity, we get what could, in honor of Carnap, be called an "extended pair list". Some of the relations of phenomenal similarity require ordered triplets, ordered quadruples, and so on, so the list is "extended" in a second sense beyond what, in Carnap, were nothing but pairs.

Theorists of many persuasions can agree that the extended pair list comprises data that must somehow or other be explained. This corpus of relational statements is compatible with sense-data accounts and with their contraries. It is compatible with those who want to be property realists about whatever sensible features are in question, and with those who deny that doctrine. It is compatible with a thorough-going intentionalism or representationalism about sensing red, and it is compatible with anti-intentionalist or anti-representationalist views.

III

So how *does* one explain looks? Let us consider how the question is approached in experimental psychology (see Koffka 1935, 75-87). Why does this portion of the scene (this stimulus) x presented to subject S in circumstances C look the same as (or bear such-and-such a relation of

phenomenal similarity to) this other portion of the scene y presented to S in circumstances C' ?

Experimental psychologists have invented various ingenious models in response. The general principle is to account for as much of the variation as possible by appeal to observable physical similarities in the objects perceived and in the circumstances of perception. The variation that remains must be explained by appeal to theoretically postulated processes within the subject who is doing the perceiving. What makes two metamers metamers, for example, cannot be specified without making reference to perceptual mechanisms within the subject. Two bundles of wavelengths that match for Jack might not match for Jill; two bundles that match for Jack in circumstances C might not match for him in circumstances C' .

Often, then, we must trace the origins of phenomenal similarities and dissimilarities inwards, into commonalities and divergences of internal states found within the perceptual processing of the subject. These models postulate internal "stages" of processing and "ensembles" of internal states linked by those processes. For example, in a classic "black box" model, each box represents some psychological process that takes certain input signals and produces an output. The boxes are connected by lines (channels) so that the output of one becomes the input of others. Each such input line or output line can vary over a range of values; that range of possible values is (in information theory) called an "ensemble" (see Shannon 1948).

The signals in an ensemble can typically be categorized or ordered. In early versions of the opponent-process color appearance model, for example, there is some stage at which the system registers the degree of yellowness/blueness of a stimulus; another process registers hue variations along an axis from red to green (Hurvich, 1981). The values are often ranges of contraries, which are mutually exclusive and jointly exhaustive. The links between stages can be unidirectional or not, and branching or not. Whatever the exact architecture, a prototypical explanation for a phenomenal similarity in respect of F is to find some stage d in which (a) when stimuli x and y both look the same in respect F to any subject S in any circumstances C , the model predicts that both x and y will cause states with the same value at stage d ; and (b) encoding in the same value at stage d is determinative of indiscriminability with respect to F thereafter. What (a) means is that two stimuli that "look the same" are shown to have the same effects within the given ensemble of internal states. What (b) means is that once they have the same effect at that stage, the two stimuli cannot thereafter be discriminated in respect of F ; they will be indiscriminable, or at least match.

To the extent that how things look cannot be explained by observable physical properties of the objects in question or of the circumstances of perception, the only recourse is to explain them by appeal to similarities and differences of such *theoretically postulated* internal states of the subject. Perhaps the properties of these states, upon which such explanations hang, *are* one variety of what philosophers have called the "qualitative" character of

sensory mental states.

IV

That last suggestion is not a new idea: it comes from Sellars (1963). Sellars argued that the very idea of an "impression" derives from an act of theoretical postulation. The concept of an "impression" is *not*, according to Sellars, derived from facts that would be obvious to any clear-headed introspector prior to the invention of those models. Instead the models provide the concepts in terms of which introspection of these internal states becomes, for the first time, possible.

Consider a particular juncture in the myth of Jones (Sellars 1963, 194-195) which is obscure, if not problematic. This juncture occurs after the natives have learned a Rylean language to make perceptual reports. Jones teaches them how to make a reporting use of the vocabulary he invented in his theoretical postulates. The problem: why should we think that this particular theoretical vocabulary (unlike others, such as Newtonian physics, string theory, or quantum chromodynamics) will be found to be learnable, and useful, for making *introspective* reports? How could theoretical postulation provide the concepts needed for the natives to report on their own mental states? As a first step in contrasting philosophical and psychological approaches to explaining looks, I will argue that the psychological models just described provide a relatively simple answer to this question. They show how these particular postulated properties of theoretical entities *could* come to be useful in making introspective reports.

These models are models of the similarities and differences in the appearances of parts of the *world*. What matches, or resembles, or fails to match or resemble, are always portions of the world--parts of the scene, portions of the ambient array--perceived by particular subjects in particular perceptual circumstances. The observable physical properties of those stimuli and those circumstances help to make some sense of the patterns of similarities and difference, but once they are exhausted these models attempt to account for the remaining variance by postulating variations in the internal states of the perceiving subject.

It follows that if one holds those first parameters constant, then according to these models variations in the appearances of things indicate variations in the internal states of perceiving subjects. For example, the fact that green things look more vividly green after one has adapted to a strong red light does not imply anything about the physical properties of the green things (which remain unchanged) or about the physical properties of other elements in the circumstances of perception (which are likewise unchanged); instead such a report on "how things appear" is indicating something about variations in the internal states of one's visual apparatus, in particular its state of adaptation.

Once some clever hominid learns this, reporting that "green things look brighter now", which has the overt form of a report about appearances of portions of the world, can come to serve as an indicator of the reporter's visual

state of adaptation. It is not that Jones invents a new, distinct reporting use of the vocabulary of appearances out of nothing; rather, according to the model itself, that vocabulary has all along tacitly served as an indicator of variations in the perceiver's internal states. (At least it does so in the circumstances in question, when all the other parameters of variation can be discounted.) So when Jones teaches the natives how to make a reporting use of his theoretical vocabulary of impressions, he does not need to invent *ex nihilo* a whole new use of the vocabulary to describe a newly constituted set of facts. Instead, all he need do is make what was tacit overt. He need only reveal that, in these particular uses of the vocabulary of appearance, used apparently to describe the appearances of things, the language-users have (in a sense) been "introspecting" all along. Variations in some of those usages are indicators of variations in their own perceptual states of mind. They simply lacked the concepts, and the theoretical models, that would enable them to see those variations as variations in their own mental states. All along, the variations in how things look turn out to reveal variations in the observers' own internal perceptual states.

A similar revelation of something that in fact was always present can be seen in Chisholm, where the innocent witness of the cabin on Mt. Monadnock is led eventually to realize that part of his evidence for his knowledge of the latter lies in his own states of being appeared blue-to. If "it looks blue to me" is true, then "I sense blue with respect to it" and "I am aware of blue with respect to it" are also true. And the latter forms are overtly saying something about one's "mental state".

Instead of casting the report in a form describing the appearances of things, Jones substitutes a form overtly about the internal state of the perceiver. The natives are taught to say, in exactly those same circumstances, something like "With me it is as if I were seeing a blue rectangular wafer". That which they notice, which prompts them to make "introspective" reports about their own "minds", might be identical to what previously prompted them to report on how things look.

V

Our "extended" pair lists record facts of the form: some portion x of the world perceived by subject S in circumstances C stands in some relation of phenomenal similarity or dissimilarity to some portion y of the world perceived by subject S in circumstances C' . And suppose in the fullness of time we acquire a stock of psychological models that explain some instances of these relations. Some of the instances can largely be explained by physical characteristics of the stimuli in question or of elements in the circumstances in question. Others though require independent sources of variation, located within the perceptual mechanisms of the perceiving subjects. Sometimes we can identify a locus--a crux within a particular stage of processing--in which the sorting into ensembles at that stage can explain why the stimuli causing the processing are perceived as similar or dissimilar. If we succeed we gain the

capacity to identify an ensemble of mental states whose occurrences and properties are determinative of relations of phenomenal similarity and dissimilarity.

There are many contrasts between such psychological modeling and the philosophical traditions for understanding appearance. I will describe five.

(1) As mentioned at the beginning, one ancient philosophical worry derives from the contrast between illusion and reality. The olive oil "looks green" in energy-saving light in that in those circumstances it looks the same as something that is green. By most common sense standards the olive oil would be judged *not* to be green; in most circumstances, it *looks* yellowy-amber, not green. To use a somewhat more modern terminology: if one has an visual experience of something that looks green but in fact is not green, then the experience is "non-veridical". Whereas if one sees something that looks green, and it is green, then the experience is "veridical". What is it for a visual experience to be veridical? How are we to understand the contents of visual experience, and the correctness conditions for those contents, in such a way as to make sense of ordinary distinctions between illusion and veridical perception? And how can we tell, in any particular episode, whether it is an episode of veridical perception or of illusion?

Concerned as it is with the grounds needed for true judgements, knowledge, etc., the locus of philosophical worry became how to *distinguish* episodes of illusion from those of veridical perception. To make that distinction the philosopher must provide some account of perceptual content, of the correctness conditions for that content, and of how we might pick out those unfortunate episodes in which a non-veridical content is presented.

It is surprising, and liberating, to realize that the relational forms can be employed and understood, and they can successfully describe a large array of facts to be explained, in complete independence of solutions to these philosophical worries. For the psychologist the critical fact is that stimuli *x* and *y* match one another; the further description of one of the episodes as "illusory" and the other as "veridical" is, from the psychologist's point of view, optional. The first task for the psychologist (and it is rather difficult one!) is to explain why the appearance of one *matches* the other. Color scientists can pursue this job while maintaining a rather insouciant indifference to the question of which, if any, of those appearances is veridical.

But mustn't a color scientist understand what it is for something to be green before they can explain to us what it is for something to appear to be green? In a word, no. The simplest way to sidestep this demand follows the lead of Cohen (2004, 2007). Sometimes the olive oil looks yellowy-amber; sometimes it looks green. If (as recommended for other reasons already) we insist on the fully relativized form, there is no actual contrariety between these various appearances. We can allow both that the olive oil looks green (to subject *S*, viewing it under the spectrum of an energy-saving fluorescent light) and that it also looks yellow-amber (in daylight, say). There is no need to identify one of these appearances as veridical, and the other as illusory. Cohen

suggests that fully relativized, they are *all* veridical, in the sense that all of them have equal claim to that label. There is no distinctive reason to say this one is veridical and the rest are not. Notice that on the same grounds they all also have an equal claim to being non-veridical. Spectators from afar, not engaged in this dispute, could reasonably conclude that here the very distinction between "veridical" and "non-veridical" loses its point. Invoking it will not lead to world peace.

(2) A similar divergence in interests between philosophers and psychologists arises over the question of what constitutes the qualitative "character" (or "content") of some mental states. Philosophical accounts of that content, and of the correctness conditions for it, have become increasingly sophisticated and diverse in recent years (see Siegel 2005). While the psychologists do not need to settle these philosophical disputes in order to proceed with their model building, their models do have certain implications about the minimal conditions that any theory of qualitative character must satisfy. The relational form makes these implications manifest.

Suppose we aim to explain why some stimulus x looks the same in hue as some other stimulus y , and the physical characteristics of those portions of the world and of the conditions of viewing cannot account for the phenomenal similarity. The model-builders posit some sequence of stages of processing, and a particular ensemble somewhere within that sequence which is critical to the hue discriminations in question, in the sense that if information about the distinctness of hue attributes of stimuli is lost at or before that stage, then thereafter the stimuli can never be discriminated in respect of hue. The way stimuli are "encoded" at that stage (the value in the ensemble which the viewing of the stimulus causes) is determinative of matching and discriminability with respect to hue. Then the psychologist aims to demonstrate that both x and y cause, at that stage, the same internal state Q_i . If the derivation is successful (if given the physical specifications of the stimuli and the details of the stages of processing in the model, we can show that by stage Q in the model both of the physically very different stimuli x and y will cause the same state Q_i), then information about the distinctness of x and y with respect to hue is forever lost, and the two will be found to match.

It follows that the Q states in this model have a simple kind of qualitative "content": one that is a species of information content. Q_i is one value in an ensemble within a sequence of information channels. The Q states are found somewhere in a channel whose inputs are stimuli. They also provide inputs to a channel whose outputs ultimately yield all the discriminations for all those relations of phenomenal similarity of which the given ensemble is determinative.

In particular, if x occasions at the Q stage the value Q_i and this stage is determinative of hue matching, then the occurrence of Q_i must carry sufficient information to ensure that x will match any member of a particular set of stimuli S_i (all those that it matches in hue), and that it will fail to match all the members of other sets S_j (those from which it can be discriminated in respect

of hue). This is true as well for all the other values Q_i that might occur in that ensemble. The variations within this ensemble must carry sufficient information to manage all the discriminations of which the ensemble is determinative.

Now whatever else "qualitative character" turns out to be, sameness and difference of the qualitative character of mental states must cohere in the right way with the holding or not of the various relations of phenomenal similarity amongst the stimuli causing those various mental states. If, to continue this example, we are trying to explain the (extended) pair lists associated with "matches in hue", then the qualitative character of the mental states involved must line up in the right way with the subject's ability to match or discriminate different hues in different *things*. If occasions yield mental states with exactly the same qualitative character, then the subject should find the two stimuli to match. If the subject can readily discriminate hue differences between two stimuli, then those stimuli had better not cause mental states that have identical qualitative character with respect to hue. Sameness in qualitative character must yield the same relations of phenomenal similarity among the stimuli those mental states encode. This is a kind of "minimal" or "fundamental" qualitative character that all the more intricate varieties must sustain.

Suppose for example we think of states of V4 as subserving a particular kind of qualitative character: they provide the "neural core" for appearances as of a hue. (Similarly one might think of the FFA as providing the neural basis for appearances as of faces.) Now the "appearance as of a hue" cannot be utterly generic; it has to be the appearance of a hue that is at least somewhat distinguishable from other hues. How are we to understand the particularities of this particular appearance? Or, in other words, what is the particular "content" of this particular appearance as of a hue? Whatever else it is, it has to be an appearance as of a hue that is contrary to at least some others, and so distinguishable, and discriminable, from some others. And if the internal states in question do indeed provide the neural core for such appearances, then their occurrence must be determinative of the discriminations in question. Given normal background conditions, what goes on in V4 must be sufficient to sustain the appearance as of a distinguishable hue, and to distinguish it from those other hues.

It follows that the qualitative character of a mental state caused by a stimulus in a given stage of processing must carry information sufficient to specify some relations of phenomenal similarity--of matching or discriminability, or whatever relations are managed at that stage--between the stimulus it registers and other stimuli. Whatever else it is, the qualitative character of that mental state suffices to specify some relations of phenomenal similarity between the stimulus that causes it and other stimuli. If we think of specifying relations to a bunch of relata as specifying a structural property, then whatever else it does, the qualitative character of a particular mental state in one of these ensembles must suffice to specify a structural property. Which structural property, and which relations, depends entirely on which relations of

phenomenal similarity the states in this ensemble are postulated to explain.

It seems fair to describe the content derived from these relational properties as a "core" or "root" notion of qualitative character. It is the trunk of the tree, if you will. There are many branch points beyond it in the philosophical literature detailing the typology of possible kinds of qualitative character. These states might be representational or not. If representational, they might be conceptual representations, or not. If conceptual, they might be Fregean or Russellian. But whatever the details, all these branches share this common, root, or core set of conditions on qualitative character. Sameness and difference of qualitative character among mental states must cohere with relations of phenomenal similarity and dissimilarity among stimuli.

(3) The contrast between representationalist models of qualitative character and non-representationalist models is worth exploring a bit further. Cast in terms of psychological modeling of sensory appearance, it can be put as follows. Does every such model include at least one ensemble of states that are representational--states that, in an older terminology, have an intentional content, and are about some intentional object? Or are there models that can explain these similarities and dissimilarities yet entirely avoid positing internal states that are representational?

This issue has recently become known in philosophical circles as "intentionalism" or "representationalism" (see Byrne 2001). The notion is that if the light in the kitchen causes my visual system to enter a state in which the olive oil looks green to me, then the only way, or the best way, to understand this is that the light in the kitchen has caused me to represent the olive oil as being green. This representation is (here) non-veridical; the experience is an illusion. Nevertheless the experience has a content, and the content is not satisfied. We explain the illusion by explaining how inputs into the visual system in those circumstances cause it to represent the color of the olive oil incorrectly. "Looks *F*" is analyzed as "is visually represented to be *F*", and this can happen even if no present stimulus is *F*. So it gives a simple yet powerful account of perceptual illusions.

Nevertheless there is no good a priori reason to expect every possible psychological model of this domain to have that form; other explanatory models might be equally successful. In fact there exists an old rival (defended by Sellars 1963) that denies intentionality--or at least dramatically restricts the attributes of intentionality--for the mental states involved. Instead of an ensemble or ensembles of states that represent, it posits internal states that have a distinctive set of mental properties, whose patterns of similarities and differences are isomorphic to the similarities and differences in the sensible properties of things one perceives. Instead of red and green (etc.) which sensibly characterize things one sees, these would be red* and green* (etc.), which are properties of internal states caused by the things one sees. These are non-intentional properties of one's sensory states; they do not strictly speaking represent anything, but they help to explain the properties that seem to characterize portions of the scene. Green* is not green; the olive oil neither

is, nor appears to be, green*. Instead the portion of the scene causes a mental state that is green*, and in virtue of causing that mental state, it looks green. Green* might be a property that distinguishes a particular Q_i state from other mental states in that ensemble.

Call this a "Sellarsian" or "homomorphism" model (see Rosenthal 2005). Obviously the dispute between the representational and non-representational models is ill-defined until one gives a precise distinction between mental states that are representational and those that are not. Or rather: until one lays out the family of distinctions governing this territory. Why is not the having of a green* visual state the representing of something being green? Contrariwise, if the content of internal visual states is cited only in the context of explaining instances of phenomenal similarity and dissimilarity, has the theorist citing them thereby invoked the apparatus of intentionality? Clearly these questions cannot fruitfully be addressed unless and until proponents give detailed accounts of exactly what the representational systems are, and why it is appropriate to invoke intentional discourse to describe them.

(4) One recent and burgeoning debate in philosophy of perception--the debate over "disjunctivism"--can be situated in terms of the legitimate *scope* of relations of phenomenal similarity. The most restrictive scope is to allow relations of phenomenal similarity only between portions of the world perceived in the "good cases": cases in which something is F and is perceived to be F . If something appears to be F , and it is not F , then (on many accounts) we have a case of "mere" appearance, or illusion, and the case is not a "good" case. Theorists who take the most restrictive line on the scope of relations of phenomenal similarity might argue that no relation of phenomenal similarity obtains between things perceived in the good cases and things that appear in any illusion. Such a view is in some ways close to the variety of disjunctivism proposed by Hinton (1973) and Martin (2004).

A less restrictive reading is to allow that things that look F or appear to be F in an illusion can look just the same as something that is perceived to be F in a "good case", but that the scope of relations of phenomenal similarity does not extend to hallucinations. Schematically, in a hallucination it seems as if one is perceiving something that is F , but in fact there exists no thing that is being perceived as being F . A common psychological description is: perception in the absence of a stimulus. This less restrictive reading allows phenomenal similarities between "good cases" and cases of illusion, but not between either of those cases and cases of hallucination. This thesis resembles the variety of disjunctivism proposed by Snowdon (1980).

The least restrictive is to allow relations of phenomenal similarity to obtain between all three cases: perceptual, illusory, and hallucinatory. Sellars (1963, 151) provides a useful example. He thought there was a "similarity among the experiences of"

- (a) Seeing that x , over there, is red
- (b) Its looking to one that x , over there, is red
- (c) Its looking to one as though there were a red object over there

If we consider "the idea that *x*, over there, is red", then this idea is endorsed in (a), "partially endorsed" in (b), and not endorsed at all in (c). From a later description it is clear that (a) is the good, perceptual case, (b) is the illusory case, and (c) is an hallucination. Sellars says there is a "similarity among the experiences" of

seeing that an object over there is red, its looking to one that an object over there is red (when in point of fact it is not red) and its looking to one as though there were a red object over there (when in fact there is nothing over there at all. (Sellars 1963, 175)

So it appears that Sellars thought that the scope of relations of phenomenal similarity was unrestricted, and could obtain between any of the cases of types (a), (b) or (c).

A Snowdon-style disjunctivist will allow phenomenal similarities to obtain between cases of types (a) and (b), but will deny them between (c) and either (a) or (b). A Hinton/Martin style disjunctivist will allow phenomenal similarities to obtain only among the (a) cases, and deny any content common between (a) cases and either (b) or (c).

Here there exists neither time nor space to describe the arguments for these philosophical positions (see Byrne & Logue, forthcoming). However, the readers of this issue might be interested to find images of this debate, within experimental psychology, cast in terms of the scope of explanatory models. It is conceivable that one kind of explanatory model will suffice for cases of types (a) and (b)--perception and illusion--but fail to explain cases (c)--hallucination. Such would be the image, within empiricism, of Snowdon-style disjunctivism. Hallucinations, since they arise without a stimulus, might involve very different mechanisms than do either perceptual experiences or illusions. Both (a) and (b) are prompted by stimuli, while (c) is not.

Likewise, it is conceivable that models that suffice to explain perceptual experience cannot be extended even to episodes of illusion, so that the scope of models explaining (a) does not extend to (b). Such would be the empirical image of a Hinton/Martin style disjunctivism. This last possibility is indeed conceivable, but the bulk of psychological theorizing about illusion assumes the contrary: that illusions involve the very same mechanisms and systems that are involved in perception. Indeed there is even some evidence that hallucinations involve the same mechanisms as the "good", perceptual cases (see Ffytche et. al 1998)

(5). Last but not least, we can finally broach the most contentious contrast between philosophical and psychological models of appearance. Many philosophers assume that phenomenal properties, the qualitative character of some mental states, qualia, etc. are coeval with, require, or (in some way) essentially involve consciousness. We even have the label "phenomenal consciousness" noted at the beginning, which suggests a rather intimate tie between the two. But is there any essential link between appearances and consciousness?

The relational form enables us to ask formulate this question precisely. Suppose we confront an extended pair list of instances of phenomenal

similarities and dissimilarities. In one instance some portion of the world x presented to subject S in circumstances C looks just the same as some other portion of the world y presented in other circumstances C' to the same subject S . To explain this and other instances of phenomenal similarities psychologists create models in which the assemblage, processing, and sorting of internal mental states explains why this m -tuple and that m -tuple (of stimuli presented on occasions to subjects) match, and these others do not. Some of the mental states involved might be "conscious mental states", in various senses. They might be mental states of which the subject is conscious, however that latter fact is ultimately analyzed or explained. They might be mental states *that* the subject is conscious of being in. Or, the weakest sense, they might be mental states the having of which makes the subject conscious of whatever object, entity, or event that subject is at the moment perceiving. Now the question. Must explanatory models of phenomenal similarities and dissimilarities posit some ensemble of states (some set of states at some stage in the system) that are, in *any* of these senses, conscious mental states?

I believe that the answer to *that* question is, quite clearly, "no". The simplest argument points to examples of models that explain a particular instance of phenomenal similarity even though none of the states involved are, in any way, conscious mental states. A good place to search for states that satisfy the latter condition is in early (or pre-attentive) visual processing. In particular it is useful to study the ways information is provided to selective attention prior to its selecting its next target for attentional processing. If those mechanisms are prone to illusions--if for example they can be shown to treat subjective contours very much like real contours--then their states would sustain relations of phenomenal similarity, even though those states are in no sense conscious states. I have argued elsewhere (Clark 2006) that there are many viable examples of such phenomena. Davis and Driver (1998) provide a particularly compelling case.

If such arguments are sound, then mental states that have qualitative character can occur without their subjects being aware of them, being conscious that they have them, or being aware of that of which they are perceptions. They can occur, in short, without awareness anywhere in the vicinity. Their occurrence nevertheless helps explain why things in the world manifest the patterns of phenomenal similarity and dissimilarity that they do. So it seems one can be appeared-to without being conscious-of. To put it more strongly, adding "awareness", in any of the three varieties just mentioned, is not necessary in these explanations. It is not the case that awareness is always needed.

VI.

As I hope I made clear at the beginning, there are many different notions of "phenomenal property" and "qualitative character" found in the marketplace of ideas, and this paper has focused on just one, tractable, kind. To close it might be useful briefly to examine one other, less tractable notion of what it is for a

mental state to possess qualitative character. This kind will be a kind that essentially involves consciousness. Recent prominent proponents of rather sophisticated versions of this idea have been Thomas Nagel (1979, 1986) and John Searle (1997, 2004), but I think a simple version can be found in the second of Thomas Reid's *Essays on the Intellectual Powers of Man* (Reid 1855). To distinguish this kind from what I take to be a tractable kind, I will follow Nagel and call this kind "phenomenological properties" as opposed to "phenomenal properties".

The root idea is that sensory appearance is essentially a subjective phenomenon, where by "subjective" what is meant is that the reality of what is happening is somehow constituted by how things appear to a subject. That reality cannot be understood therefore without grasping how it appears to that subject. Furthermore, subjective phenomena are claimed to have a different metaphysical or ontological status than other, "objective" phenomena (see Nagel 1979, 213; Nagel 1986; Searle 1997, 98, 114; Searle 2004, 135). One version is that if a state such as having a pain is a "subjective" phenomenon, or (equivalently) if it involves phenomenological properties, then if it appears to a subject *S* that *S* is in that state, then *S* is in that state. That state has phenomenological property *F* if and only if it appears to *S* that it does. The appearance of it constitutes the reality of it. Whereas this is not so for quotidian "objective" properties such as size or shape. A second version of the difference in metaphysical status between "subjective" and "objective" phenomena is that the former can only be understood if one can understand the point of view of the subject on the phenomenon (Nagel 1979, 188). Whereas understanding objective phenomena requires no graspings of or takings of points of view.

An interpretation of some of the passages in Thomas Reid's second *Essay* can yield a simple version of this notion of phenomenological properties. The two key claims are that there exists a set of mental states ("sensations"), such that (a) if such a state appears to its subject to be *F*, then and only then is it *F*; and (b) the having of such states requires the subject to be conscious of them. In fact, perhaps it is because the subject is conscious of them that the reality of such states is constituted by their appearance (as claimed in (a)). In the essay Reid never explicitly commits himself to these two theses about sensations, but a number of passages have been interpreted by subsequent readers in a way that would commit Reid to those claims. The consequence is that Reid (or a possible Reid) has been thought to hold the view that sensations are mental states of a rather special kind.

The critical passages are meant to explain the notion that sensations have "no objects distinct from themselves". Reid explains this as follows:

This sensation can be nothing else than it is felt to be. Its very essence consists in being felt; and when it is not felt, it is not. There is no difference between the sensation and the feeling of it; they are one and the same thing. It is for this reason that we before observed, that in sensation there is no object distinct from that act of the mind by which it is felt; and this holds true with regard to all sensations. (Reid 1855, 141)

In another passage he says

Sensation is a name given by philosophers to an act of mind, which may be distinguished from all others by this, that *it has no object distinct from itself*. Pain of every kind is an uneasy sensation. When I am pained, I cannot say that the pain I feel is one thing, and that my feeling it is another thing. They are one and the same thing, and cannot be disjoined even in imagination. Pain, when it is not felt, has no existence. It can be neither greater or less in degree or duration, nor any thing else in kind, than it is felt to be. It cannot exist by itself, nor in any subject but a sentient being. No quality of an inanimate, insentient being can have the least resemblance to it. (Reid 1855, 140)

Now there is a reading of the claim that "There is no difference between the sensation and the feeling of it; they are one and the same thing" which implies both (a) and (b) above, and I believe this interpretation has been historically potent. First, we get (a), because a sensation is what it is felt to be. If it feels as if it is *F*, then it is *F*. Furthermore, it is *nothing* else than what it is felt to be. So if it is not felt to be *F*, then it is not *F*. The conjunction of those two claims yields (a). To get (b) we need to note two features of the very complicated English verb "to feel". One has already been applied: it is a verb of appearance. "Feels cottony" is analogous to "looks green", except the former describes the appearance of something one is touching instead of seeing. But "feel" also has a sense of "sensible of" or "aware of", and the form "feeling *F*" generally implies one is aware of *F*. So Rosenthal (1997, 732) argues, I think successfully, that "feeling the headache" implies "being aware of the headache", whereas just "having a headache" might not.

In the passages just cited it is clear that the "feeling of the sensation" is an act of mind, requiring a sentient being, in virtue of which that sentient being is aware of that sensation. Its "feeling the pain" entails that it is conscious of the pain. And this is (b).

Now (a) and (b) conjoined make "sensations" a remarkable class of mental states, and their properties, which I dubbed "phenomenological properties", would be remarkable too. They manifest a simple form of the weirdness found in Nagel's "subjective facts". First, it is a remarkable state of mind that could be such that it is exactly and only as it appears to be. It cannot have properties it does not appear to have, and whatever properties it appears to have it must have. Sensations thereby become existents for which the appearance/reality distinction collapses. For them whatever features seem to be real are real. How could there be such existents? What reason do we have for supposing that there are any such things?

The answer to that question draws on the second claim about these phenomenological properties: that they must be conscious. Ordinary states of affairs, of the sorts for example that one perceives, are such as to sustain some distinction between appearance and reality. The reality can contradict the appearances, and, sometimes shockingly, include properties that are not at all apparent. But perhaps, the interpretation goes, states of affairs inside consciousness have a different metaphysics. That they are inside consciousness is another way of saying: the ordinary rules do not apply.

Because they are states of consciousness, there is no difference between the sensation and the feeling of the sensation. Only special states, within a sentient being, need apply. This idea attaches a special metaphysical status (a capacity to generate reality out of mere appearance) to the state of being conscious.

Historically one sees the implications of this interpretation working themselves out in the claims that sensations *must* be conscious, and that consciousness is a different sort of ontological domain, to which the ordinary rules do not apply. Both ideas have done significant damage over the history of psychology.

I suggest the proper response for an empiricist is to ask for evidence that states with this character exist at all. What is the evidence that there exist states of mind that have all and only the properties that they appear to their subject to have? It is hard to conceive a circumstance in which an experimental psychologist could reasonably posit theoretical entities that have these properties. "At this point in the model there arise elements x such that each x has all and only the properties that it appears to have." A caption for a cartoon, perhaps, but unimaginable in serious psychology. But if the demand for such metaphysically potent entities does not arise from the need to explain observations, from whence does it arise? Why should we credit this suggestion at all? What reason is there to think that such properties and entities are even possible?

References

- Barwise Jon. (1981). Scenes and other situations. *Journal of Philosophy* 78 (7): 369-397.
- Barwise, Jon and Perry, John (1983). *Situations and Attitudes*. Cambridge, MA: MIT Press.
- Broad, Charles Dunbar (1927). *Scientific Thought*. New York: Harcourt Brace & Company Inc. Originally published London: Routledge and Kegan Paul, 1923.
- Byrne, Alex (2001). Intentionalism defended. *Philosophical Review* 110 (2): 199-240.
- Byrne, Alex and Logue, Heather (forthcoming). Either/or: disjunctivism for dummies. In A. Haddock and F. McPherson (eds.), *Disjunctivism: Perception, Action, Knowledge*. Oxford: Oxford University Press.
- Carnap, Rudolf (1967). *The Logical Structure of the World*, 2nd edn., tr. Rolf A. George. Berkeley: University of California Press.
- Chisholm, Roderick (1957). *Perceiving: A Philosophical Study*. Ithaca, New York: Cornell University Press.

- Clark, Austen (2006). Perception Preattentive and Phenomenal. In Paul Thagard (ed.), *Handbook of Philosophy of Psychology and Cognitive Science*, Dordrecht: Elsevier, pp. 105-139.
- Cohen, Jonathan (2004). Color properties and color ascriptions: A relationalist manifesto. *Philosophical Review* 113(4): 451-506.
- Cohen, Jonathan (2007). A relationalist's guide to error about color perception. *Nous* 41(2): 335-353.
- Davis, Greg and Driver, Jon (1998). Kanizsa subjective figures can act as occluding surfaces at parallel stages of visual search. *Journal of Experimental Psychology: Human Perception and Performance* 24(1): 169-184.
- Ffytche, Dominic H., Howard, R. J., Brammer, M. J., David, A., Woodruff, P., and Williams, S. (1998). The anatomy of conscious vision: an fMRI study of visual hallucinations. *Nature Neuroscience* 1(8): 738-742.
- Hinton, John M. (1973). *Experiences: An Inquiry into Some Ambiguities*. Oxford: Clarendon Press.
- Hurvich, Leo M. (1981). *Color Vision*. Sunderland, Massachusetts: Sinauer Associates Inc.
- Koffka, Kurt (1935). *Principles of Gestalt Psychology*. New York: Harcourt Brace and World.
- Martin, Michael G. F. (2004). The limits of self-awareness. *Philosophical Studies* 120: 37-89.
- Nagel, Thomas (1979). Subjective and objective. In his *Mortal Questions*. Cambridge: Cambridge University Press, 196-214.
- Nagel, Thomas (1986). *The View from Nowhere*. Oxford: Oxford University Press.
- Reid, Thomas (1855). *Essays on the Intellectual Powers of Man*. Edited by James Walker, D. D.. Sixth Edition. Boston: Phillips Sampson and Company.
- Rosenthal, David (1997). A theory of consciousness. In Ned Block, Owen Flanagan, and Güven Güzeldere (eds), *The Nature of Consciousness*. Cambridge, MA: MIT Press, 729-754.
- Rosenthal, David M. (2005). Sensory quality and the relocation story. In his *Consciousness and Mind*. Oxford: Clarendon Press, 149-174. (Originally published in *Philosophical Topics* 26(2), 1999, 321-350.)
- Searle, John R. (1997). *The Mystery of Consciousness*. New York: New York Review of Books.
- Searle, John R. (2004). *Mind: A Brief Introduction*. Oxford: Oxford

University Press.

Sellars, Wilfrid (1963). Empiricism and the philosophy of mind. In his *Science, Perception, and Reality*. London: Routledge, 127-196.

Shannon, Claude E. (1948). A mathematical theory of communication. *Bell System Technical Journal* 27: 379-423, 623-656.

Siegel, Susanna (2005). The contents of perception. In *The Stanford Encyclopedia of Philosophy*. Edited by E. N. Zalta. URL = <http://plato.stanford.edu/entries/perception-contents/>

Snowdon, Paul (1980). Perception, vision, and causation. *Proceedings of the Aristotelian Society* 81: 175-92.