

Version of September 23, 2006

Visual Awareness and Visual Qualia*

“It is of the essence of mental paint to be something of which we are aware.” Ned Block

“No awareness without representation!” Anonymous sage

Christopher S. Hill
Department of Philosophy
Brown University
Providence, RI 02915

Jaegwon Kim has done as much as anyone to clarify the relationship between mental phenomena and the phenomena that are studied in the physical sciences. On the whole, his account of the relationship provides encouragement and support for the view that the world is fundamentally unified – that its mental dimension is intimately connected with its physical dimension, and is, moreover, connected with it in a way that is intelligible through and through. Early and late, however, Kim has felt that it may not be possible to accommodate qualitative mental states, or qualia, within his picture of an integrated and unified world. In the end, he thinks, it may be necessary to see qualia as an irreducible “mental residue.”¹

I will try here to continue Kim’s integrative efforts by arguing for an account of awareness of qualia that promises to bring qualia within the physical fold. More particularly, I will present and defend a version of the view that has come to be called “representationalism.” I will focus on the task of developing a representationalist theory of visual awareness, but it will be evident, I think, that the theory can be generalized so as to apply to experiential awareness of other kinds as well. After explaining this theory of visual awareness, I will urge that it provides a satisfactory answer to the metaphysical problems that qualia pose. I will then discuss the question of where exactly qualia are to be located within our catalog of physical properties. In the final sections of the paper I will attempt to answer the most pressing objection to representationist accounts of visual awareness – the inverted spectrum objection.

I

Although there are differences of opinion as to how many types of qualia there are, there is wide agreement that qualia are associated with bodily sensations, emotions, and

perceptual experiences. Thus, one category of qualia includes such characteristics as *being a pain, being an itch, and being a surge of nausea*. Another category consists of the properties like the ones that we invoke when we speak of the icy hand of fear and the glowing coals of anger. And a third category consists of properties that are associated with the ways that objects appear to us when we perceive them. As Kim says, perceptual qualia “are, by definition, the ways that things *look, seem, and appear* to conscious observers.”²

In order to characterize this third category of qualia adequately it is necessary to distinguish between two senses of the terms that we use to describe appearances. I will focus here on “looks” since I will be concerned primarily with visual qualia in the present paper, but my remarks will apply, *mutatis mutandis*, to “seems,” “appears,” and a number of other appearance words.

There is a sense of “looks small” in which it can be correctly applied both to a toy car that one holds in one’s hand and to a real car that one sees on the road far ahead. In this sense, the expression can also be applied to a tall building that one sees from a plane, and even to an immense star that one sees from the Earth. When one says that an object looks small to an observer, using “looks small” in this *phenomenological* sense, one is not claiming that the observer’s perceptual experience supports the judgment that the object really is small. One is not saying that the observer’s experience represents the object as small. Rather one is drawing an analogy between the observer’s current visual experience and the visual experiences he has when is viewing objects that are reasonably close at hand and really are small.

The phenomenological sense of “looks” is also to be found in claims about apparent shape and apparent color. It is permissible to apply “looks elliptical” both to an object that really is elliptical and is perpendicular to the observer’s line of sight, and to a round coin that is slanted away from the observer. Equally, it is permissible to apply “looks dark brown” both to a piece of chocolate and to a portion of a tan wall that is cloaked in shadow.

In addition to the phenomenological sense of “looks,” there is also what is often called its *epistemic* sense. When we say that an item looks small to an observer, using “looks small” in this second sense, we mean that the observer’s current visual experience provides adequate evidential support for the belief that the object is small. When we have this second sense in mind, we would not be willing to say that a car looks small to an observer if the car is at an appreciable distance from the observer, for when a car is at an appreciable distance from an observer, the observer’s visual experience presents him with “pictorial cues” that are indicative of distance. Thus, for example, when an object stands at some distance from an observer, the features of the object seem indistinct. In a case of this sort, the observer’s experience supports the belief that he is seeing a car of normal size, but a car that is rather far away. Accordingly, using “looks” in its epistemic sense, it is correct to say that the car looks to be of normal size, despite the fact that it is also correct to say, using “looks” in its phenomenological sense, that the car looks small to the observer.

We can also use “looks” in its epistemic sense to talk about appearances of other kinds. Thus, it is quite appropriate to apply “looks round” to a coin that is tilted away from an observer, and to apply “looks tan” to a portion of a wall that is poorly

illuminated, provided that the observer's visual experience attests to this fact about the lighting.

Putting the distinction between these two senses of "looks" to use, we can define visual qualia as characteristics that we are aware of in virtue of the ways that objects look_p to us. (Here and hereafter, I use "looks_p" to represent the phenomenological sense of "looks.") In general, perceptual qualia can be defined as the characteristics we are aware of in virtue of the ways in which objects appear to us, where "appear" is used in its phenomenological sense.

I have mentioned three categories of qualia – sensory qualia, emotional qualia, and perceptual qualia. What do the members of these categories have in common? Why do we group them together under a common label? One reason is that they seem to us to share a number of epistemological properties. We believe that it is impossible to know qualia adequately without experiencing them. We also believe that our grasp of them is direct and immediate. Moreover, we are strongly inclined to think that our awareness of qualia is not governed by an appearance/reality distinction. Thus, we are inclined to think, it is impossible for it to seem to one experientially that one is in pain without one's actually being in pain. Equally, we suppose, it is impossible for it to seem to one that an object looks_p red to one without its actually being the case that an object looks_p red to one. Finally, we are inclined to think that experiential awareness provides us with full access to the essential nature of qualia. Our grasp of them is not perspectival or limited in any way. They do not have a hidden dimension that experience fails to reveal.

These epistemological perceptions would by themselves provide a sufficient reason for viewing qualia as special, but we are also strongly inclined to suppose that qualia

have unusual metaphysical properties. It is probably our metaphysical perceptions that lead us to think of them as qualitative, and to use “qualia” as a term for them.

We view qualia as intrinsic properties – as properties that things have independently of their relations to other things. Also, we are inclined to think that there is something metaphysically fundamental about qualia: when qualia are not themselves are simple and unanalyzable, we are inclined to suppose, they are resolvable into more basic qualia that *are* simple and unanalyzable. Further, we hold that qualia are responsible for relationships of qualitative similarity and qualitative difference. This last claim is explained in different ways by different authors, but it seems to be common ground that qualitative similarity is a form of similarity that cannot be analyzed in terms of shared causal powers or shared spatio-temporal relations. It is a form of similarity that arises directly from the intrinsic natures of objects. Moreover, since the qualia that are responsible for it are simple and unanalyzable, qualitative similarity is simple and unanalyzable as well. In particular, qualitative similarity is not analyzable in terms of shared microphysical structure.

Here, then, is a list of properties that we conceive of qualia as possessing. But do qualia really have the properties on this list? Or, to rephrase the question, are there characteristics that count as qualitative in the sense that they possess the properties on the list?

Certainly it is hard to locate characteristics with the listed properties in the physical world. Consider, for example, the characteristic we are aware of when something looks_p yellow to us. From the perspective of experience, we are strongly inclined to say that this characteristic is simple, in the sense that it cannot be resolved analytically into other

characteristics that are in some sense more basic.² But the physical characteristics that figure in scientific accounts of color vision are all complex. Science attributes complexity to the enduring physical characteristics of the external substances that serve as the objects of perception, to the viewpoint dependent characteristics that objects possess in virtue of reflecting structured light to various vantage points, and to the characteristics that various parts of the visual system come to exemplify during the various stages of visual processing.

Further, when we grasp qualia experientially, they seem to us to have individual natures that are distinct from the individual natures of all of the characteristics that science investigates. Consider again the characteristic that we are aware of when something looks_p yellow to us. Call it *C*. Are we aware of *C* when we visually explore the brain, or when we visually imagine the microprocesses in the brain that science describes? No. Are we aware of it when we examine the retina, perhaps aided by a powerful microscope? No. Are we aware of it when we visualize in imagination the interactions of light waves and the surfaces of external objects? No. Are we aware of it when we visualize the enduring spectral reflectances of such objects? No. In view of facts of this sort, we are strongly inclined to think that there is no experience we can have of physical characteristics, or can imagine ourselves having, that would put us in touch with *C*.

Now if there were an appearance/reality distinction associated with visual qualia, we could simply say that we have trouble locating *C* in the physical world because the appearance it presents to us fails to reveal its true nature. That is, we could explain away the apparent difference between *C* and all physical characteristics. Unfortunately,

however, it seems inappropriate to distinguish between C as it appears to us and C as it is in itself. C is the ostensibly qualitative characteristic that is affiliated with facts of the form $x \text{ looks}_p \text{ yellow to } y$. Since there is no appearance/reality distinction that is associated with these facts, how could there be an appearance/reality distinction that is associated with C ?

In general, it seems, when an object looks_p a certain way to us, we are aware of a characteristic that we cannot be aware of in any other way. No other form of experience will reveal it to us. Moreover, when we experience it, we are not aware of it as having a microphysical nature of any sort, or as determining relationships of similarity of the sort that arise from microphysical structure. Accordingly, we sense that it is different than the characteristics that science describes. And we are unable to explain this sense of difference away by invoking an appearance/reality distinction.

Alas, these reflections seem to lead inexorably to the conclusion that there is an unbridgeable gulf between qualia and the physical world.

Here then is the crux of the problem that qualia present to us. We feel obliged to embrace a form of qualia realism. After all, it seems that we are aware of qualia. But there is nothing in the physical world that answers to our conception of qualia. Hence, qualia cannot be physical characteristics. The physical world does not exhaust reality. In Kim's apt phrase, there is a "mental residue."

II

As far as I can tell, there is only one theory of experience that provides even a remotely satisfactory way of dealing with this problem – representationalism.

Representationalism maintains that our awareness of the characteristics we call qualia

essentially involves representations, and that the representations in question are different in a variety of respects than the representations that are involved in other forms of cognition. Because of the distinctive features of these representations, it maintains, the properties they represent seem to us to have special features like intrinsicness and simplicity, and seem to us to have individual natures that are not captured by scientific accounts of experience. Despite our impressions to the contrary, representationalism asserts, our awareness of qualia is governed by an appearance/reality distinction of a certain sort. There is the way that the properties we call qualia are in themselves, and the way that these properties seem to us in virtue of being represented by the representations that subserve experiential awareness. It is only insofar as they are represented by such representations that they take on the aspect of qualia.

Representationalism maintains, in other words, that the properties we call “qualia” do not really have the properties that we take to be constitutive of qualitative character. They only seem to us to have such features because of the peculiarities of the representations in virtue of which we are aware of them. Moreover, if they seem to us to have individual natures that are different from the natures of all physical characteristics, this is because, and only because, they are represented in a unique way. Representationalism adds that we are unable to see beyond these apparent differences, and to appreciate the ultimate identity of the properties we call “qualia” with certain physical properties, because it is not apparent to us, from the perspective of common sense, that our awareness of them involves representations. Folk psychology does not reveal that representations are constitutively involved in facts of the form *x looks_p F to y*. Accordingly, it does not occur to us that our awareness of the properties we call “qualia” might be governed by an

appearance/reality distinction. We think that it is necessary to take our experience of qualia at face value.

This view of qualia can also be expressed, albeit somewhat opaquely, by saying that qualia are not properties that exist independently of our awareness of them. They are properties as seen from the perspective of the systems of representation that enable experiential awareness. *They are physical properties qua experientially represented.*

In developing this story, representationalism begins by calling attention to some of the ways in which experiential awareness differs from the forms of awareness that involve conceptualization and judgment. Among other things, it points out that experiential awareness always provides us with access to highly determinate forms of properties (when a thing looks_p blue to one, it always looks_p a highly determinate shade of blue), that experiential awareness generally provides more information about objects of awareness than can easily be captured by linguistic or conceptual descriptions, that it automatically assigns the properties it represents to locations in quality spaces, that it enables us to make extremely fine discriminations among the members of domains, thereby giving us a sense of the density of the orderings that obtain in the domains, and that it is associated with a variety of attention mechanisms – mechanisms that enable us to adjust certain aspects of experiential awareness, such as resolution and figure/ground contrast.

According to representationalism, it is possible to explain these and other distinguishing features of experiential awareness by invoking features of the systems of representation that such awareness involves. Of course, at present this has the status of a largely unsubstantiated hypothesis. We know very little about the representational

systems that are involved in experiential awareness. We have but a meager grasp of their formal properties, and views about their semantic properties are largely conjectural. Fortunately, we do not need to have a well worked out account of experiential representation in order to see that it is in principle *possible* to provide representation-based accounts of the foregoing features. Suppose, for example, that the systems of representation that are involved in experiential awareness are analog in character. On this assumption, we can easily visualize a representation-based explanation of why experiential awareness always puts us in touch with highly determinate forms of the properties it represents, and also an explanation of why it generally provides us with access to the density of the orderings that govern the domains with which it is concerned. The mere possibility of giving explanations of this sort is sufficient to make representationalism an attractive theory. Advocates of representationalism at least have the right to *hope* that that the distinguishing features of experiential awareness can be explained. It is not clear that any other theory of experiential awareness confers this right.

Thus far we have been considering the general motivation for representationalism. I will try now to say a bit more about the relevance of representationalism to the metaphysical problem that arises when we suppose that there really are characteristics answering to our conception of qualia. In these remarks, and also in later parts of the discussion, I will focus on visual qualia – that is, on the qualia we are aware of in virtue of participating in facts of the form *x looks_p F to y*. I believe that what I will say about visual qualia generalizes to qualia of other types, but I will not attempt to defend this belief here.

Our awareness of visual qualia is experiential in nature. Now as we saw a bit earlier, there is independent motivation for supposing that there is a distinctive system of representation that subserves experiential awareness. In view of this fact, representationalism contends, we have the right to assume that our awareness of visual qualia is essentially representational, and that the representations involved in such awareness belong to a distinctive system.

Applying this assumption, we can explain why certain properties seem to us to be simple and unanalyzable, despite having an internal complexity of the sort that science reveals. Thus, it is a familiar fact that a representation can encode information about a property without encoding information about the details of its internal organization. We know, for example, that a swath of green paint on a canvass can indicate the presence of vegetation on a hillside without indicating whether the vegetation is grass, or wheat, or ivy, or kudzu, or some other sort of ground cover. A fortiori, it doesn't tell us anything about the physical microstructure of the vegetation. Nor does it tell us anything about the physical nature of greenness. It represents the hillside as covered with green vegetation, period.⁴

Further, there is a representationalist explanation of why there seems to be an unbridgeable gulf between the property we are aware of when something looks_p yellow to us and the properties that are revealed by the scientific investigation of vision.

When things look_p yellow to us, we are deploying a representation that is different than all of the representations that we deploy when things look_p other ways to us, different than all of the representations that are involved in non-visual experiential awareness, and different than all of the conceptually structured representations that

science makes available. Let us say that this representation represents the characteristic *C*. Now of course, the mere fact that we use a special representation to keep track of *C* could not by itself give rise to an abiding impression that we are aware of a characteristic that is distinct from all other characteristics when we are aware of *C*. This merely creates the possibility of such an impression. The impression itself arises from our sense that it would be inappropriate to identify *C* with a characteristic that we grasp via some other representation.

We are all familiar with the fact that there can be multiple representations of a single characteristic, and we all avail ourselves, from time to time, of the option of identifying a characteristic that we initially grasp via one representation with a characteristic that we initially grasp via a different representation. But when we identify one characteristic with another, we are obliged to appeal to the fact that the characteristics are represented by distinct representations in explaining and justifying the identification. Only in this way is it possible to account for the fact that it is necessary to go beyond what experience itself reveals about the characteristics in order to appreciate the real nature of their relationship.

Now we must ask: Is it appropriate to invoke the difference between a property-qua-represented and a property-as-it-is-itself in determining whether *C* is identical with some other property? From the perspective of common sense, the answer is “no.” We are aware of *C* in virtue of participating in facts of the form *x looks_p yellow to y*; and when we view such facts from the perspective of common sense, we find no reason to suppose that they constitutively involve representations. That is to say, folk psychology affords no glimpse of the representations that are constitutively involved in our experience of *C*. We can only appreciate the existence of those representations from a highly theoretical

perspective. (Representationalism is a highly theoretical position.) Accordingly, we do not see how it could be true that *C* might have a nature that is not revealed by experiential awareness. Since there is no apparent distinction between *C*-as-represented-by-experience and *C*-as-it-is-in-itself, it cannot possibly be appropriate to go beyond what experience reveals about *C* in assessing *C*'s relationship to other characteristics.

It may be useful to recast this line of thought in a somewhat different form. Suppose that *X* is a characteristic that is revealed by experience, and suppose also that for some reason it is deemed desirable to identify *X* with a *prima facie* different characteristic *Y*. In order to explain and justify the identification, which involves going beyond the impression of *X* that experience itself provides, it is necessary to explain how it is possible to grasp *X* experientially without appreciating its identity with *Y*. It is normally possible to provide such an explanation by invoking some sort of appearance/reality distinction. Thus, we might distinguish between *X* itself and a property that serves as the mode of presentation for *X*. Or, if there is no other property that serves as the mode of presentation for *X*, as will be the case if our awareness of *X* is direct, then we might distinguish between *X*-as-it-is-in-itself and *X*-as-it-is-represented by an experiential representation. But, to repeat, we must invoke some such contrast in order to explain why the identity of *X* with *Y* is not revealed by experience itself.

Now let us turn to consider the special case of the characteristic *C*. Is it possible to identify *C* with some other characteristic, say *C**? If we are to do so, there must be a way of explaining how it is possible to grasp *C* experientially without appreciating its identity with *C**. This means that we must invoke an appearance/reality distinction of some sort. But folk psychology does not recognize a distinction between appearance and reality in

this case. It fails to register the representational character of our awareness of *C*, and by the same token, it fails to support any ambitions that we might have to identify *C* with another characteristic.

This completes my account of how representationalism explains our impression that *C* is distinct from all other characteristics. Of course, in addition to explaining that impression, it offers a perspective from which it is appropriate to reject the impression as illusory. Unlike folk psychology, representationalism affirms the representational nature of our awareness of *C*.

Thus far I have made six claims. It may be useful to summarize them. First, I have claimed that when an object *x* looks_p *F* to an observer *y*, *y* is aware of *x* as having a certain property, a property that is invoked by the locution “looks_p *F*.” I will henceforth speak of this form of awareness as *experiential awareness*, and I will say that the properties that are objects of experiential awareness, the properties that are invoked by predicates of the form “looks_p *F*,” as *appearance properties*. Second, I have claimed that experiential awareness is representational in character. It constitutively involves a representation of an appearance property. Third, I have claimed that experiential representations have distinctive properties that set them apart from representations of other sorts. Thus, for example, unlike conceptual representations, experiential representations are analog or quasi-analog in character. Fourth, I have claimed that visual qualia are appearance properties, and that experiential awareness is the form of awareness that puts us in touch with qualia. Fifth, I have claimed that it is possible to explain the special metaphysical properties that qualia seem to us to have, such as simplicity, intrinsicness, and position in a quality order, in terms of the special nature of experiential

representations. Thus, for example, it is possible to explain why it seems to us that qualia are simple by appealing to the fact that experiential representations of qualia do not represent them *as* complex – that is, by the fact that such representations do not encode information about their internal organization. Finally, I have claimed that qualia are ultimately physical in character. They can be located within our independently motivated catalog of physical properties.

It is clear that the task of explaining and defending these claims is quite large. In the present paper I will be concerned only with two of them – the claim that awareness of qualia is essentially representational in character, and the claim that qualia can be found within the space of physical properties. These are perhaps the most important of the foregoing claims, but the others are important too. I hope to consider the others on a future occasion.

III

There are three main alternatives to representationalist theories of experiential awareness. One alternative maintains that experiential awareness is a matter of *direct acquaintance* with facts involving objects and certain special properties. On this view, experiential awareness is simple and unstructured. It is a primitive cognitive relation that can obtain between observers and certain facts. It occupies the ground floor with respect to explanation and analysis. All that can be said of it is that it is a form of awareness, and that the properties that are given to us by this form of awareness are qualitative. The other alternative to representationalist theories is *adverbialism*. Roughly speaking, adverbialism claims that when an object looks_p *F* to an observer *y*, what is going on is that *y* is aware of *x* in a certain *way*. To be a bit more specific, according to adverbialism,

y is aware of x *F-ishly*. Thus, while adverbialism allows that qualia exist, it maintains that they exist only as *forms of perceptual awareness* – as *ways* of being perceptually aware of non-qualitative phenomena. That is to say, adverbialism maintains that qualia are adverbial qualifications of an underlying generic relation of perceptual awareness, a relation that agents bear to external objects. The third alternative to representational theories of experiential awareness is what might be called the *doxastic theory*. According to this view, experiential is propositional and doxastic in character, and therefore necessarily involves some sort of conceptualization. To be aware of a quale is to token a representation of a property, but the representation in question is fully conceptual in character. All knowledge of qualia takes the form of judgments.

I think there are strong reasons for preferring representationalist accounts of awareness to accounts that are based on acquaintance. Thus, in the first place, representationalist accounts give us some hope of being able to account for the ways that objects of awareness appear to us. In particular, what is especially relevant to our concerns, they afford some hope of our being able to account for the ways that the properties that we call “qualia” appear to us. These properties seem to us to be intrinsic, seem to be simple, or at least to be analyzable into simple components, seem to support relations of qualitative similarity and qualitative difference, and so on. As noted earlier, representationalist theories of awareness have some promise of explaining why certain objects of awareness seem to us to have these characteristics.

In addition to providing a foundation for explaining our impressions concerning the objects of awareness, representationalist theories provide a basis for answering certain key questions about the nature of awareness itself. At any given time one is aware of

certain objects and characteristics and not of others. Why is one aware of the items that one is in fact aware of? Why isn't one aware of the others instead? Further, why is the scope of awareness broad at some times and not at others? Why does the resolution of awareness change when it does? What is the difference between attentive awareness and more casual forms of awareness? We are not yet in a position to give detailed answers to questions of this sort by appealing to properties of the system of representation that supports experiential awareness, but elementary reflection shows that it is possible in principle to answer them in this way. Thus, for example, it is clearly possible to explain differences in the scope of awareness by appealing to differences in the contents of representations: the more inclusive the content of the representations that are involved in awareness, the broader the scope of awareness. It is also possible to explain differences in resolution in terms of differences in the contents of representations. But theories that invoke acquaintance in their accounts of awareness are ipso facto precluded from answering questions about how awareness works. Acquaintance is supposed to be simple and fundamental. If this is true, then facts involving acquaintance are in an important sense brute facts. They cannot be reductively explained.

In view of these considerations, it is clear that representationalist theories of experiential awareness are overwhelmingly superior to those based on acquaintance in point of explanatory power. This is a good reason for preferring them. Another, related reason is that they provide us with a simpler picture of the mind. They enable reductive explanations of otherwise recalcitrant facts, thereby substantially reducing our ontological commitments and streamlining our account of nomological relationships.

I turn now briefly to adverbial theories of experiential awareness. According to these theories, it will be remembered, qualia exist only as *forms of perceptual awareness* – as *ways* of being perceptually aware of non-qualitative phenomena.

I see no merit in this suggestion. Adverbialism implies that there is no such thing as perceptual awareness of qualia. Now it would be grotesquely implausible for a theory to maintain that there is no awareness of qualia *of any kind*. We are not simply ignorant of qualia, as we would be if there were no such thing as awareness of them. Clearly, for example, there is such a thing as awareness of pain. When one is aware of a pain, one is aware of a quale. Thus, if adverbialism wishes to deny that we are perceptually aware of qualia, it must identify some other form of awareness and claim that awareness of qualia involves *that* form of awareness. But that other form can't be adverbial, on pain of an infinite regress. At the end of the day, therefore, adverbialism must claim that we are aware of qualia in one of the ways we have already distinguished – via direct acquaintance, or via judgments, or via experiential representations. Adverbialism adds complexity without adding substance.

This brings us to the doxastic theory of experiential awareness. This view allows that experiential awareness constitutively involves representations, but it differs from the view I am defending under the label “representationalism” in that it claims that the representations in question are conceptual through and through. It maintains that to be aware of a quale is to make a conceptually informed judgment of a certain sort.⁵

We have already taken note of some facts that call the doxastic theory into question. Thus, as we observed earlier, experiential awareness has all of the following properties:

- (a) It provides us with access to highly determinate forms of properties. (When an

- object looks_p blue to me, it always looks_p a highly determinate shade of blue.)
- (b) It provides more information about objects of awareness than can easily be captured by linguistic or conceptual descriptions.
 - (c) It automatically assigns the properties it represents to locations in quality spaces.
 - (d) It enables us to make extremely fine discriminations among the members of domains, thereby giving us a sense of the density of the orderings that obtain in the domains.
 - (e) It is associated with a variety of attention mechanisms, including mechanisms that enable us to adjust certain aspects of experiential awareness, such as resolution and figure/ground contrast.

Awareness that involves conceptualization and judgment lacks all of these features. It is not true, for example, that conceptual awareness always provides us with access to highly determinate forms of properties. I can judge that something is blue without judging that it has a highly determinate shade of blue. Nor does conceptual awareness have property (b) – indeed, it would be contradictory to say that it does. It cannot be true that conceptual awareness provides us with information that cannot be captured in conceptual terms. Nor does conceptual awareness have property (c). A blind person can have the capacity to form judgments about various shades of color without knowing how to locate those shades within the color solid. Nor does it have property (d). Anyone who has experiential access to colors will appreciate that similarity orderings of shades are dense, or are at any rate very finely graduated. On the other hand, a blind person, even if he has somehow managed to acquire an immense number of color concepts, will not automatically have a sense of how finely graduated such orderings are. Nor does conceptual awareness have

property (e). To be sure, we do speak of attention in connection with conceptualization and judgment. For example, it is possible to attend in thought to a theorem in number theory. But this sort of attention is different in kind than the forms of attention that are involved in experiential awareness. Thus, in the case of doxastic attention, there is nothing that is strictly analogous to increasing resolution and figure/ground contrast.

This list of differences could easily be extended.⁶ Experiential awareness must be distinguished from doxastic awareness.

We have now found strong reasons for setting all of the main alternatives to representationalist accounts of experiential awareness aside. I will henceforth assume that some sort of representationalist account is correct. More specifically, I will assume that it is correct to say that awareness of qualia constitutively involves representations of a sort that can appropriately be called *experiential*.

IV

I have been concerned thus far only with the nature of awareness of qualia. Before going on to consider the nature of qualia themselves, I would like to emphasize the point that a representationalist theory of awareness of qualia is sufficient by itself to solve the metaphysical problem of qualia that we considered at the outset. Or at least, a representational theory is sufficient provided that it can be developed in a satisfactory way. However desirable it may be to supplement the theory with a positive account of the nature of qualia, it is not necessary to do so in order to block the main arguments for property dualism, and to bridge the gulf separating qualia from the physical world that we considered at the outset.

The metaphysical problem of qualia derives from the fact that qualia seem to us to have properties that all of their physical correlates lack. Thus, it seems to us that they are simple and intrinsic, and that they occupy positions in such quality spaces as the color solid. The physical properties that are correlated with qualia lack these characteristics. It appears that representational theories of experiential awareness have the capacity to solve this problem. A representational theory makes it possible to draw an appearance/reality distinction with respect to qualia. On the one hand, there are qualia-as-they-are-represented-by-experiential-representations. On the other hand, there are qualia-as-they-are-in-themselves. Because it is possible to draw this distinction, we are not obliged to say that qualia really are simple and intrinsic, nor that they really do occupy positions in quality spaces like the color solid. Moreover, it is plausible that, given an appropriate account of the nature of experiential representations (an account which claims, among other things, that such representations are analog in character), it will be possible to give a detailed explanation of why qualia seem to us to have these characteristics. It will be possible to close the explanatory gap. My present point is that it is possible to draw a distinction between the appearance of qualia and the corresponding reality, and to develop a theory of experiential representations that explains the appearances of qualia, without having a theory of the nature of qualia. *The mind-body problem arises because folk psychology provides us with an incomplete and otherwise faulty account of qualitative awareness. Accordingly, it can be dealt with by developing a better theory of awareness. It is not necessary to supplement such a theory with an account that explains exactly what kind of properties qualia are.* Relative to the mind-body problem, any positive theory of the nature of qualia is a luxury.

V

This is not to deny that it is independently desirable to understand the nature of qualia. It is clear that this is an important goal. I turn now to the task of constructing such an account. Thus far, all that has been said about visual qualia is that they are appearance properties – or in other words, that they are the properties we are aware of in virtue of participating in facts of the form $x \text{ looks}_p F \text{ to } y$. The goal now is to consider the issues that arise when one attempts to go beyond this starting point.

There are four views about appearance properties that seem, at the outset, anyway, to be genuine options:

First view: Appearance properties are properties of internal mental entities of some sort, such as sensations or sense data.

Second view: Appearance properties are objective, physical properties of external objects – properties like objective shape and objective size.

Third view: Appearance properties are causal properties of the form *being an external cause of an internal mental occurrence with intrinsic property Q*.

Fourth view: Appearance properties are viewpoint dependent physical properties of external objects such as *subtending a visual angle of A degrees, undergoing angular displacement at rate R, reflecting light of absolute intensity I in direction D, and reflecting light of spectral composition C in direction D*. That is, they are properties that external objects have in virtue of their relations to *vantage points*, where vantage points are not internal to observers, but are rather positions in physical space that observers can occupy.

These are not the only views about the nature of appearance properties that have appeared in the history of thought, but they are certainly among the most prominent. It is plausible, initially at least, that one of them must be correct. Unfortunately, as I will try to show, they are weighty objections to all of them. None of them, it seems, can be true – or at least, none of them can be said to provide the whole truth and nothing but the truth.

I will be concerned here only to evaluate these traditionally important views. I will not go on to present and defend a fifth view. The task of finding an appropriate alternative to the four traditional accounts of appearance properties seems to me to be one of the hardest, and also one of the most important, of the problems facing the philosophy of perception. It is also a problem for vision science. My goal in this paper is just to explain some of the issues that must be addressed in any search for a new and better view.

Could the first view be correct? That is, could appearance properties be properties of internal, mental entities of some sort? I think we can see that the answer is “no” by reflecting on the transparency of visual awareness. Visual awareness is transparent in the sense that the only objects that are presented to us in visual awareness are external, physical objects. We are not aware of internal objects of any kind. Now this implies that any properties we are aware of in visual experience must be properties of external objects. Since we are not aware of any internal objects, we could not be aware of any properties of internal objects unless we were aware of them without being aware of them as characterizing objects, that is, as instantiated. It seems absurd to say that visual experience presents us with properties that are floating free of all objects, as mere possibilities of instantiation. Thus, transparency provides us with a reason to say that appearance properties, and therefore qualia, are properties of external objects.

Could appearance properties be objective, physical properties of external objects – that is, properties like objective shape and objective size? It seems that the answer must be “no.” Appearance properties are properties that we are aware of in virtue of participating in facts of the form $x \text{ looks}_p F \text{ to } y$. Such facts are essentially perspectival. The appearances that objects present to us are constantly changing, but the objective, physical properties of objects are relatively constant. In view of this fact, it seems wrong to identify appearance properties with objective, physical properties.

This leaves us with two possibilities. One is that appearance properties are properties of a sort that I will call *Galilean*. They are properties of the form *being an external cause of an internal mental occurrence with intrinsic property Q*. The other possibility is that appearance properties are what I called above *viewpoint dependent* properties of external objects – that is, properties that objects have in virtue of relations to physically determined perspectives or vantage points, properties such as visual angle and the spectral composition of reflected light. I will now sketch an argument which suggests that the Galilean option is more problematic than the viewpoint option. Then I will mention a consideration that seems to call the viewpoint option into question. In both cases, I will focus on what the theories say about apparent *size*

I begin by observing that there is a law of mathematics that makes it possible to recover the objective size of an object from its visual angle, given that one knows the objective distance of the object, and that also enables one to recover the objective distance of an object from its visual angle, given that one knows its objective size.⁷ On the other hand, there is no law of mathematics, nor even a law of nature, that makes it possible to recover objective size or objective distance from a Galilean causal property.

This is because there is no law that makes it possible to recover objective size and objective distance from an intrinsic property of internal mental occurrences. What distinguishes any one Galilean property from all others is just the intrinsic mental property that it involves. Because Galilean properties depend for their individuation in this way on intrinsic properties of mental occurrences, the fact that there are no lawful relations linking intrinsic properties of mental occurrences to objective size and objective distance implies that there can be no such relations linking Galilean properties to these other properties.

Next, I wish to claim that if information of one type T1 is quite valuable to the organism, while information of another type T2 is of little or no value, then it is more likely that there are states of the visual system that have the function of encoding information of type T1 than that there are states that have the function of encoding information of type T2. I hope that this principle seems plausible. I will not argue for it here.

Now it is quite clear, in view of the fact that visual angles are related lawfully to properties that are of great importance to the organism, that it is potentially quite valuable for the organism to have information about visual angles. Also, it is correspondingly clear that the organism has little to gain from possessing information about Galilean properties. Accordingly, it is more plausible that the visual system has states that have the function of encoding information about visual angles than that it has states that have the function of encoding information about Galilean properties. Following Dretske, we can express this by saying that it is more plausible that the visual system has states that *represent* the former properties than that it has states that *represent* the latter properties.⁸

We have found, in effect, that the representations that support awareness of apparent sizes are more plausibly regarded as representing visual angles than as representing Galilean properties. This leads naturally to the view that apparent sizes should be identified with visual angles. Moreover, it is clear that the argument can be generalized, permitting us to draw the general conclusion that it is more plausible that appearance properties should be identified with viewpoint dependent properties than that they should be identified with Galilean properties.

Although there are no laws linking intrinsic properties of mental occurrences to the objective properties that are of primary importance to the organism, it is plausible that there are laws linking *informational* properties of mental occurrences to objective properties of this sort. Accordingly, an advocate of Galilean properties might propose that the foregoing simple formulation of the Galilean view should be replaced with the thesis that appearance properties are causal properties of the form *being an external cause of an internal mental occurrence with informational property I*. The trouble with this suggestion is that it appears to be parasitic on the view that the mind makes use of information about viewpoint dependent properties. If it is useful for the visual system to represent a property of the form *being an external cause of an internal mental occurrence with informational property I*, this must be because *I* is information about a property that is related lawfully to objective physical properties, for it is objective physical properties that are of primary importance to the organism. Now as we have seen, there are viewpoint dependent properties that have simple lawful relationships with objective physical properties – relationships that can easily be exploited by computational procedures. Moreover, it is hard to see what other properties there might be that enjoy

such relationships. Accordingly, if it is useful for the visual system to represent a property of the form *being an external cause of an internal mental occurrence with informational property I*, it is plausible that the information in question is information about a viewpoint dependent property.

It appears, then, that the revised Galilean view implies that the mind makes use of information about viewpoint dependent properties. In view of this fact, it seems best to set the Galilean view aside. Whereas the Galilean view implies that the mind represents causal properties that incorporate information about visual angles, the alternative view claims that the mind represents visual angles. It is clear that the latter view is much simpler and much more straightforward than the former.

There are other ways of revising the foregoing simple version of the Galilean view, but perhaps enough has been said already to show that the hypothesis that the visual system represents visual angles and other viewpoint dependent properties enjoys a prima facie advantage over the hypothesis that it represents causal properties that involve characteristics of internal mental occurrences. The burden of proof is on the Galilean's shoulders.⁹

Unfortunately, the view that apparent sizes are visual angles faces a serious problem. Roughly speaking, the problem is that there are considerations which suggest that apparent sizes are not strictly proportional to visual angles. Rather, it seems, the apparent size of an object reflects the results of early *constancy computations* concerning objective size. That is to say, the apparent size of an object seems to be a kind of weighted average of visual angle and an early estimate of objective size. This fact has been known to vision scientists for some time. It has also been known that similar things can be said of

apparent shape, apparent color, apparent motion, and so on. Here, for example, is a passage concerning apparent shape from an influential paper that appeared in 1931:

If a subject is shown an inclined circle and is asked to select from a number of figures the one which represents the shape seen by him, he chooses without hesitation an ellipse. This ellipse, however, is widely different from the one which represents the shape of the inclined circle indicated by the laws of perspective, being much nearer to the circular form. The subject sees an inclined figure neither in its 'real' shape nor in the shape which is its perspective projection but as a compromise between them.¹⁰

The author of this passage, Robert Thouless, makes similar observations concerning appearance properties of several other sorts. Thouless's conclusions are sustained by modern research.

If apparent sizes are not strictly proportional to visual angles, then they cannot possibly be identical with visual angles. But we must be careful here. The apparent size of an object x is the property one is aware of in virtue of participating in a fact of the form $x \text{ looks}_p F \text{ to } y$. It is the property that is represented by the experiential representation that is constitutively involved in the fact. It could be true that this property is a visual angle even though the agent makes judgments concerning apparent size which imply that apparent size is not proportional to visual angle. Thus, it could be the case that the relevant representation *represents* the visual angle that is subtended by x , but that it *distorts* or *misrepresents* the size of that angle. We know that there is such a thing as systematic misrepresentation in perception, though we do not have a satisfactory philosophical account of it.¹¹ Perhaps that is what is happening here. In other words, it

could be the case that apparent sizes are actually identical with visual angles, but that they *appear* not to be proportional to visual angles because our experience misrepresents their true nature. As soon as we allow that representations figure in facts of the form x *looks_p F to y*, we must also recognize the possibility that such facts involve distortions. This possibility could be relevant to the present problem.

It turns out, then, that the view that appearance properties are identical with viewpoint dependent properties faces a problem. The problem may admit of a solution, but it is clear that the task of evaluating all of the relevant considerations is both complex and delicate. It would be a mistake to rush to judgment here.¹²

To summarize: There are problems facing all of the traditional views concerning the nature of appearance properties. At present, there is no clear answer to questions about the metaphysical nature of these properties – nor, by the same token, to questions about the metaphysical nature of visual qualia.

VI

In earlier sections I was concerned to explain a particular form of representationalism, to argue that it promises to provide a satisfactory solution to the metaphysical problem of qualia, and to maintain that this version of representationalism is *prima facie* superior to competing accounts of visual experience. Insofar as these efforts were successful, we have reason to think that this version of representationalism is correct.

At this point I would like to turn to consider a difficulty -- the inverted spectrum problem. This is by no means the only problem that representationalism must contend with, but in the judgment of many, it is the most serious problem. Along with others, Kim thinks that it is insurmountable.¹³ Also, it is perfectly general. It confronts all

representationalist proposals, not just one that I have been concerned to present and defend in the foregoing discussion.

The inverted spectrum objection can be seen as an argument with two premises, the first of which is the following inversion thesis:

(IT) It is in principle possible for there to be two observers, y and z , two objects, x and w , and a form of color experience *looking_p F* such that (i) x looks_p F to y , (ii) w looks_p F to z , and yet (iii) the representations that y uses to keep track of the color of x have different contents than all of representations that z uses to keep track of the color of w .

The second premise is the following thesis about visual qualia:

(QT) If x looks_p the same way to y as w looks_p to z , then the qualia that y experiences in virtue of being visually aware of x are the same as the qualia that z experiences in virtue of being visually aware of w .

Combining these principles, we obtain the conclusion that it is possible for y and z to be experientially aware of the same qualia even though any color-related property that y may be representing concurrently is different than any color-related property that z may be representing concurrently. This conclusion spells trouble for representationalism, for representationalism claims that the qualia that are presented by visual experiences are properties that those experiences represent. If qualia are properties that are represented by experiences, then qualia cannot be identical when the properties that are represented by experiences are different.

My discussion of the inverted spectrum objection will proceed in two stages. I will begin by giving a quick argument to the effect that (IT) begs the question against the

representationalist. Any clear sighted representationalist, I will maintain, will think that there is a strong prima facie case against (IT). Then, later on, I will try to develop this prima facie case against (IT) into a detailed and systematic argument. My strategy will be to identify the main grounds for accepting (IT), and to show that when these grounds are examined reflectively, they can be seen to beg the question.

Here is the quick argument: In effect, (IT) claims that it is possible to hold all facts of the form $x \text{ looks}_p F \text{ to } y$ fixed (where F is a color property) while allowing the representational contents of all of the relevant visual representations to change. Now clearly, if we are to understand this claim, we must consider the nature of facts of the form $x \text{ looks}_p F \text{ to } y$. What exactly is going on when an object x looks_p F to an observer y ? When we reflect, we see that the answer to this question, or at least a big part of the answer, is captured by the following principle: If x looks_p F to y , then, necessarily, y is experientially aware of the object x , and y is also experientially aware of an appearance-property P that is possessed by x (or is otherwise associated with x). Now in combination with this principle, (IT) has the following consequence: it is possible to hold facts of the form $y \text{ is experientially aware of } x \text{ as instantiating (or as otherwise associated with) the appearance-property } P$ fixed while allowing the representational contents of all of the relevant visual representations to change. But from the point of view of a representationalist, this consequence of (IT) is absurd. Holding that representations are always constitutively involved in facts of experiential awareness, a representationalist will see the consequence as equivalent to the claim that it is possible to hold a certain range of representational facts fixed while changing the representational contents of the relevant representations, which is contradictory. Thus, from the perspective of the

representationalist, (IT) leads to absurdity when it is combined with a well motivated theory of experiential awareness. It follows that any clear sighted representationalist will reject (IT).

If (IT) can be seen so quickly to beg the question against the representationalist, why has it seemed to so many philosophers that the inverted spectrum argument poses a serious problem for representationalism? One reason is that it is by no means obvious that representations are constitutively involved in facts of the form *x looks_p F to y*. It cannot be emphasized too much that representationalism is a highly theoretical position. Its merits are not apparent from the perspective of commonsense. The representational character of facts of the form *x looks_p F to y* emerges only when one has seen and accepted the arguments that provide support for representationalist accounts of awareness.

VII

I turn now to the task of elaborating this short argument into an objection that systematically addresses the motivation for (IT). As I see it, there are two main sources of the plausibility of (IT). They are both thought experiments. I will describe one of these experiments in the present section, and will then argue that it begs the question against representationalism. I will describe the second thought experiment in section VIII. There are significant differences between it and the first experiment, but in the end, I will maintain, it involves the same sort of error. It too begs the question against representationalism.

Suppose that *x* is red, that *w* is green, and that *y* is a normal human observer. Further, suppose that *z* is an observer who belongs to a different species than *y*. The pathways

leading from z 's retinas to z 's color processing centers are inverted relative to y 's pathways, but in all other respects, z 's visual system is exactly the same as y 's. It follows from these initial assumptions that the visual processing that x occasions in y is exactly the same, at all points after the point of inversion, as the visual processing that w occasions in z . This is because any differences in the inputs that are caused by x and w are reversed at the point of inversion.

We are supposing, then, that y 's central processing of inputs from x is the same as z 's central processing of inputs from w . The next step is to draw an inference from this similarity. Thus, we now infer that w looks_p the same way to z as x looks_p to y . Observe that this inference is extremely natural. It is extremely plausible that the inference is a good one. (Call this the *supervenience intuition*. I will return to it in a moment.)

Let us now turn to consider the representations that figure in the higher level visual processing of our two observers, y and z . Since the higher levels of y 's processing are the same as the corresponding levels of z 's processing, any representations that figure in y 's processing must also figure in z 's processing, and conversely. That is to say, y 's higher level representations must be formally or syntactically the same as z 's. Let R be a higher level representation that is used in y 's visual system to encode information concerning the color red. Perhaps it is used to encode information concerning objective red, considered as a physical property of external objects, or perhaps it is used to encode an appearance property that is associated with red, such as the property of occurrently reflecting more long wave-length light than middle wave-length light. Let us just say that it is used to encode information about a certain red-related property $A1$. R is also a representation in z 's visual system. However, because z 's early wiring is inverted relative

to y's, R must have a different representational content in z's system than in y's. Let us say that it is used to encode information about a certain green-related property A2. In short, as used by y's system, R represents A1, but as used by z's system, R represents the complementary property A2. R is in this respect typical of all of the representations that y's system uses to encode red-related information. All such representations have a different representational content as used by y's system than as used by z's.

We have now arrived at the following results: (i) x looks_p a certain way to y; (ii) w looks_p the same way to z as x looks_p to y; and (iii) where R is any representation that y's system might be using to encode information that is pertinent to the color of x, R has a different content in z's visual system than it does in y's system. A case of this sort seems to be exactly what is needed to establish (IT).

I return now to the supervenience intuition. It can be expressed by the following principle:

(SP) If x causes the same higher level processing in y as w causes in z, where sameness consists in a sharing of neurophysiological properties, then x looks_p the same way to y as it does to z.

Evidently this principle is immensely appealing. But why? Why are we so strongly inclined to think that visual appearances supervene on neurophysiological properties?

The answer, I think, is that we see (SP) as following from the following two more basic principles:

(P1) If x causes the same higher level processing visual processing in y as w causes in z, where sameness consists in a sharing of neurophysiological properties, then the *visual experience* that y has in virtue of being aware of x is of the exactly the

same type as the *visual experience* that z has in virtue of being aware of w.

(P2) If the visual experience that y has in virtue of being aware of x is of the same type as the visual experience that z has in virtue of being aware of w, then x looks_p the same way to y as w does to z.

In effect, (P1) claims that experiences that occur in the course of visual perception supervene on the neurophysiological character of the attendant brain processes. (P2) claims that how an object looks_p to one is fully determined by the nature of the experience one has in virtue of being aware of the object. If we now ask why (P1) and (P2) are plausible, the answer is clear: (P1) seems to receive strong support from science, and (P2) is a more or less ground level principle, with little need of support from other quarters.

To expand a bit on the plausibility of (P1), it has been known since the work of Penfield that visual experience can be manipulated by electrical stimulation of visual centers.¹⁴ In more recent times fMRI studies have established reasonably tight correlations between many aspects of visual experience and forms of neural activity. These investigations make it overwhelmingly plausible that there are laws of some sort linking visual experiences to brain processes.

We have been considering the content of a thought experiment that is intended to support (IT), the major premise of the inverted spectrum argument, and also the probable intuitive basis for that thought experiment. Turning now to the task of evaluating the experiment, I think we can see, on reflection, that it has a disabling weakness. It is true that (P1) receives support from science. If, however, we think that visual experience is essentially representational in character, we will not be inclined to see either the

electrophysiological evidence or the fMRI evidence as establishing full strength nomological determination, but rather as showing only that there are various *local* laws connecting experiences with brain processes – that is, laws that hold only when certain boundary conditions are in place. In particular, if we think that visual experience is essentially representational, and also hold to the idea that at the most basic level, the representational content of visual experience is determined by natural selection, then we will see the scientific evidence as revealing only certain laws that link *human* visual experiences with brain processes. We will allow that imaginary creatures like our observer z, creatures who share our neurophysiology but who have a different selectional history, will be governed by different psychophysical laws, because the representational contents of their visual experiences are different than ours.

In other words, it is impossible to put (P1) forward in the service of an anti-representationalist argument without begging the question. Any clear sighted representationalist will reject it, for any clear sighted representationalist will wish to insist that visual experience essentially involves representation.

I will not repeat the foregoing reasons for thinking that visual experience is essentially representational. Suffice it to say that the hypothesis appears to be our best hope of explaining our sense that experiential awareness puts us in touch with qualitative properties. I would, however, like to make a new point – that the hypothesis is also our best hope of understanding how visual experience fits into the cascade of visual processing that vision science describes. From the perspective of vision science, visual processing can be analyzed in terms of representations and operations on representations. The output of each stage of processing is a representation that serves as an input to the

next stage. Accordingly, to see visual experience as lacking in representational content is to see it as irrelevant to visual processing. And, by the same token, to see facts of the form $x \text{ looks}_p F \text{ to } y$ as existing independently of the contents of experiences is to see such facts as isolated and sterile, as altogether lacking a *raison d'être*.

VIII

This brings us to a second, somewhat more complicated thought experiment. As was the case with the previous thought experiment, the point of this second one is to provide motivation for the inversion thesis (IT).

Let y and z be as they were in the first thought experiment, but now suppose that, at a certain time T , our human observer y undergoes a procedure that changes the wiring running from his retinas to his opponent processing center, thereby inverting his color spectrum. Thus, prior to time T , y 's wires are inverted relative to z 's, but after time T y 's visual system is structurally the same as z 's at all points.

Now as a result of this procedure, y 's post- T color discriminations are exactly the same as z 's, and the same is true of y 's dispositions to group colored together on the basis of similarity. If z would discriminate between the colors of two objects, y would do so as well. And if z were to sort a set of objects into groups, in such a way as to reflect perceived similarity of color, then y would repeat z 's sorting if he were to be given the same task. In view of these facts, it is extremely tempting to say that after time T , all objects look_p the same way to y as to z , at least as far as color is concerned. After all, the discriminations and similarity assessments one makes are determined by how things look_p to one: it is part of our folk psychological practice to explain color discriminations

and assessments of color similarity by making assumptions about the ways that objects appear to observers.

What should we say about the post-T contents of the representations that y's visual system uses to keep track of colors and color appearances? I think we must say that the inversion procedure does not change the contents of these representations. To be sure, the procedure changes the *causal* relationship between these representations and distal stimuli, and also the relationship between the representations and proximal stimuli. The representations that used to be triggered by red light are now triggered by green light, and conversely. But the *contents* of the representations that are now triggered by red light continue to be green-related, and the *contents* of the representations that are now triggered by green light continue to be red-related. We are obliged to hold this because it is extremely plausible that content in the visual system is determined by natural selection. Surgical procedures can change the causal relations between y's representations and stimuli, but these cannot change the contents of those y's representations. Representational content is independent of the vagaries of causal control.

Drawing these observations together, we find that we have a case in which all of the following seem to hold: (i) after time T, red objects look_p a certain way to y; (ii) red objects look_p the same way to z as they do to y; and (iii) the contents of y's color representations are the same after time T as they were prior to time T. Because of (iii), this is also a case in which the contents of y's color representations after time T are different than the contents of z's color representations. Thus, on the face of it, it is a case that provides strong support for the inversion thesis (IT).

The crucial step of this line of thought, I believe, is the move from the observation that *y*'s discriminations and similarity judgments coincide with *z*'s to the claim that objects look_p the same way to *y* as to *z*. We observed that this transition is plausible because we are accustomed to explaining discriminations and similarity assessments by appeal to facts of the form *x looks_p F to y*. But this is only part of the story. It is also true that from the perspective of common sense, we have no choice but to explain discriminations and similarity assessments by invoking such facts. Folk psychology provides no other form of explanation. Because of this, it seems that the conclusion that objects look_p the same way to *y* as to *z* is forced upon us.

In fact, however, even though it is extremely tempting, the inference we are considering is highly questionable. The inference would be legitimate if the claim that an object looks_p the same way to *y* as to *z* had no content beyond the claim that *y* and *z* are in a state with the same causal powers with respect to discriminations and assessments of similarity. Reflection shows, however, that when we claim that an object looks_p a certain way to an observer, we are saying much more about the observer than that he is in a state that supports certain discriminative tendencies and certain dispositions to group objects together on the basis of perceived similarity. In addition, we are attributing a specific type of experience to the observer, and therefore, a certain form of awareness. Now I have maintained above that awareness must be understood as representational in character. We are aware of objects and properties only insofar as we represent them. It simply will not do, for example, to maintain that awareness is a matter of direct acquaintance. Assuming that this is correct, if we say that an object looks_p the same way to *y* as to *z*, we are

implicitly saying that y and z are in the same representational state, and in particular, that y is in a state with the same representational content as z.

In view of these considerations, we see that it cannot after all be true that red objects look_p the same way to y as to z. This would hold only if y and z were in states with the same representational content when they were aware of red objects. In fact, however, as we have observed, even though red objects trigger representations in y's visual system that are formally or syntactically the same as the representations that they trigger in z's visual system, the representations in question have different contents in y's visual system than they do in z's system.

In effect, what we are observing here is a consequence of the descriptive inadequacy of the conceptual scheme that folk psychology makes available for describing visual experience. If we wish to account for the fact that y and z have the same dispositions to discriminate and to form similarity groups, we should appeal to the fact that the representations that red objects trigger in y are formally or syntactically the same as the representations that they trigger in z. It is the formal properties of the representations that support the dispositions. On the other hand, if we wish to say what it is that y is aware of when he views a red object, or to say what it is that z is aware of, then we should invoke the contents of the representations in question. Further, if we are to give a full account of those contents, we must say that they are different in y's case than in z's, due to the differences in the evolutionary histories of their respective species. In short, if we are to account for the facts in our admittedly rather outré case involving y and z, it is absolutely essential that we recognize the representational nature of visual experience, and further, it is essential that we distinguish sharply between the forms of the relevant representations

and their representational contents. Unfortunately, folk psychology fails to recognize the representational structure of facts of the form $x \text{ looks}_p \text{ to } y$. A fortiori, it does not distinguish between the forms of the representations that are involved in such facts and their contents. Because of these failings, it leads us to misunderstand cases like the one we have been considering, and betrays us into erroneous conclusions about them.

To conclude: it appears that there is no good reason to accept (IT). Efforts to support it make one of the following assumptions:

If y and z are in the same neurophysiological state, then objects look_p the same way to y as to z .

If y and z have the same dispositions to discriminate between objects and to form similarity groups, then objects look_p the same way to y as to z .

Both of these assumptions beg the question against representationalism. They hold in everyday cases, but they have no general validity. Accordingly, we cannot allow them to constrain our theorizing about visual experience. They should be rejected, and the inverted spectrum argument should be rejected as well.

IX

As we saw in the early going, representationalist theories of awareness appear to have the capacity to solve the metaphysical problem that qualia pose. Moreover, representationalist theories are supported by weighty considerations. If we are to obtain a satisfactory account of awareness of qualia, and of qualia themselves, it appears that we must set non-representationalist theories aside and seek an acceptable version of representationalism.

After arriving at this conclusion, we went on to consider the nature of appearance properties, and found that there are problems facing all of the traditional theories in this area. There must be a way of locating appearance properties within our independently motivated catalog of physical properties, but today it far from clear how to do this.

In more recent sections we have been considering the chief objection facing representationalism -- the inverted spectrum argument. According to this objection, it is possible to hold awareness of visual qualia fixed while changing the representational contents of all of the visual representations that could conceivably be involved in qualitative awareness. We found that the objection begs the question against representationalism. Anyone who favors a representationalist account of awareness of qualia will hold that any change in the contents of awareness-level representations will occasion a corresponding change in qualitative awareness.

In this paper I have been concerned to effect a rapprochement between the realm of matter and the realm of mind. More particularly, I have tried to show that representationalism gives us reason to think that it is possible to bring qualia into the fold of physical properties. Representationalism promises to solve the metaphysical conundrum that Kim has called the problem of "mental residue."

Notes

* I thank the Center for Consciousness at Australian National University, and its Director, David Chalmers, for support while I was working on this paper, and also for providing me with an immensely stimulating philosophical environment. I have also benefited from conversations with Uriah Kriegel and James Van Cleve, and from discussions following talks at ANU, Monash University, and the University of University.

1. Jaegwon Kim, *Physicalism, or Something Near Enough* (Princeton: Princeton University Press, 2005), pp. 170-171.

2. Jaegwon Kim, *Philosophy of Mind*, 2nd edition (Cambridge, MA: Westview Press, 2006), p. 225.

3. Here I am prescinding from a complication that is not relevant to the point at issue. Suppose that an object looks_p yellow to you, and that you are aware of a certain qualitative characteristic *C* in virtue of this fact. Here *C* is epistemically basic. But is it metaphysically basic? It is presumably possible to distinguish between three aspects of *C* – aspects corresponding to the hue, saturation, and lightness of physical colors. Each of these aspects is a property. We have a choice between saying that *C* is in some sense analyzable in terms of these three properties, and saying that they are second order properties that *C* exemplifies but that in no way constitute *C*. In saying that *C* seems to us to be metaphysically simple, or basic, I am here assuming that this second view of the relation between *C* and the other three properties is correct. But there would be no

problem about accommodating the other view of the relationship. To effect such an accommodation, I need only change my example of a simple property, citing one of the other properties in place of *C*.

4. Presumably greenness is identical with a physical property of some sort. If so, then in representing the hillside as green, the swath of green pain represents it as having that physical property. It does not, however, articulate the physical structure of the property. That is to say, it does not encode information about its structure in such a way as to make it cognitively available to the viewer.

5. In the present paper I use “representationalism” as a name for the view that awareness of qualitative properties constitutively involves *experiential* or *perceptual* representations. But there is also a more inclusive sense of the term – a sense in which it is true to say that “representationalism” stands for *any* view which claims that awareness of qualitative properties involves representations. On this more inclusive sense, it is not required that the relevant representations be experiential in character. Now beliefs and judgments are representational states, though the representations from which they are constructed are conceptual rather than experiential. Accordingly, since the third proposal claims that introspective awareness of qualitative states involves beliefs or judgments, it counts as a form of representationalism, when the term is used with its more inclusive sense. I emphasize, however, that it does *not* count as a form of representationalism when the term is used in the sense that is operative in the present paper.

6. I discuss several other features that distinguish experiential awareness from experiential awareness in “Ow! The Paradox of Pain,” in Murat Aydede (ed.), *Pain: New*

Essays on its Nature and the Methodology of its Study (Cambridge, MA: MIT Press, 2005), pp. 75-98.

7. What I have in primarily in mind here is the following law linking visual angle to objective size and objective distance:

$$\text{Height}(x) = 2 \times \text{Distance}(x) \times \tan(\text{Visual angle subtended by } (x)/2).$$

But there are other laws linking visual angle to objective properties. See the interesting discussion in Robert J. Schwartz, *Vision: Variations on Some Berkeleyian Themes* (Oxford: Blackwell, 1994).

8. See, e.g., Fred Dretske, "Misrepresentation," in Radu J. Bogdan (Ed.), *Belief* (Oxford: Oxford University Press, 1986), pp. 17-36.

9. Perhaps it will be useful to discuss one more revision of the simple Galilean view. It might be suggested that appearance properties have the form *being an external object that is disposed to cause an internal mental occurrence with intrinsic property Q when external circumstances of type E obtain*. Even if there are no laws linking simple Galilean properties to objective physical properties of external objects, it might be urged, it could still be true that there are laws of the following form: If an external object is causing an internal mental occurrence with intrinsic property Q in external circumstances of type E, then, given that such-and-such other conditions obtain, the object has objective physical property O. Perhaps the visual system exploits such laws in arriving at representations that attribute objective physical properties to external objects.

Here I will just make two quick points. First, the content *being an external object that is disposed to cause an internal mental occurrence with intrinsic property Q when external circumstances of type E obtain* is much more complex than the content *being an*

external object that subtends visual angle V. Because of this complexity, the claim that evolutionary processes have endowed the visual system with the capacity to represent the former content is much less plausible than the claim that such processes have endowed the visual system with the capacity to represent the latter content. Second, the view under consideration is committed to the further claim that we are endowed with computational procedures that incorporate or implicitly represent the laws that connect objective physical properties of external objects to properties the form *being an external object that is disposed to cause an internal mental occurrence with intrinsic property Q when external circumstances of type E obtain*. For we would only be able to recover objective physical properties from causal properties of the given form if we had an implicit grasp of such laws. Now the laws in question will inevitably be vastly more complex than the simple trigonometric principle that connects visual angle to objective distance and objective size. Hence, in addition to claiming that the visual system represents contents of considerable complexity, the present proposal is committed to claiming that it makes use of computational processes that are implausibly complex, and by the same token, it is committed to attributing an implausible complexity to the evolutionary processes that are responsible for those processes. These are additional costs of the view.

10. See Robert Thouless, "Phenomenal Regression to the Real Object I," Journal of Psychology XXII (1931), 339-359. The quoted passage occurs on p. 339.

11. Dennis Proffitt has made a strong case that there are systematic, pervasive distortions in our high level visual representations of the slants of hills and mountains. See Dennis R. Proffitt, Mukul Bhalla, Rich Gossweiler, and Jonathan Midgett, "Perceiving Geographical Slant," *Psychonomic Bulletin and Review* 2 (1995), 409-428.

There are also considerations which strongly suggest that our introspective awareness of bodily damage involves systematic distortions. See my “Ow! The Paradox of Pain,” which is cited in footnote 4. Unfortunately, while it is reasonably clear that systematic misrepresentation can occur in perception, we do not yet have a philosophical theory of representation that explains, in a satisfactory way, how systematic misrepresentations is possible.

12. One of the reasons that the issues are complex is that there are several initially promising ways of explaining the data that Thouless describes. I mentioned one way in the text – say that appearance properties are identical with viewpoint dependent properties, but maintain that the relevant experiential representations distort the nature of appearance properties, with the result that they *seem* not to be proportional to viewpoint dependent properties. Another, quite different approach is to say that appearance properties are identical with objective physical properties, but to maintain that the relevant experiential representations distort the nature of appearance properties, with the result that they *seem* not to be proportional to objective properties. That is to say, we could try to save the view that appearance properties are objective, physical properties by the same strategy that can perhaps be used to save the view that appearance properties are viewpoint dependent properties. A third possibility is to say that an apparent size is a kind of weighted average of the visual angle that is subtended by an object and an early estimate of objective size, and to say the same thing, *mutatis mutandis*, about all other appearance properties. It is a virtue of this third option that it takes our experience of appearance properties at face value. Unfortunately, the tasks of working out the details of the idea and evaluating the result will probably prove to be quite challenging.

13. Kim, *Physicalism, or Something Near Enough*, pp. 169-70. See also Kim, *Philosophy of Mind*, p. 227
14. Wilder Penfield and Theodore Rasmussen, *The Cerebral Cortex of Man: A Clinical Study of Localization of Function* (New York: Macmillan, 1950).