

Real Presence*

Alva Noë

UC Berkeley
noe@berkeley.edu

When you approach an object, it looms in your visual field. When you move around it, its profile changes. In these and many other ways, how things look depends on what you do. Competent perceivers are not surprised by these changes in appearance as they move.

Of course, objects don't usually appear to grow as we approach them; nor does it look as though they change their shape when we move. Perceptual constancy – size and shape constancy – coexists with perspectival *non*constancy. Two tomatoes, at different distances from us, may visibly differ in their apparent size even as we plainly see their sameness of size; a silver dollar may look elliptical -- when we view it from an angle, or when it is tilted in respect of us – even though it also looks, plainly, circular.

* I am grateful to Michael Martin for illuminating discussion of the ideas in this paper. For criticism and discussion, I would also like to thank James Genone, Steven Gross, William Lycan, Fiona MacPherson, Hilary Putnam, Susanna Schellenberg, James Stazicker, Michael Tye, the members of my 2004 and 2005 seminars at UC Berkeley, and Bob van Gulick and the participants of his July 2005 Syracuse Philosophy Annual Workshop and Network on consciousness.

Perceptual experience presents us with the world (the constancies) and it presents us with how the world perceptually seems to be (the nonconstancies). A satisfying account of perception must explain how the silver dollar can look both circular and elliptical, how the tomatoes can look to be the same in size and yet different in size. The content of perceptual experience is two dimensional, and this needs explaining.

Philosophers have tended to chicken out when it comes to the two-dimensional character of perceptual experience. They deny the legitimacy of one or the other dimensions of content. Sense-datum theorists deny that we visually experience the coin's circularity, or the tomato's sameness of size. At best, perhaps, we infer the presence of those features on the basis of what is present to consciousness. Direct realists, in contrast, may be tempted to deny that there is any sense in which it is *the coin* that looks elliptical, or *the tomatoes* that appear to differ in size.¹

A solution would be to hold that perception is a way of coming into contact with the way things are, apart from how they perspectively present themselves, *by coming into contact with the way they present themselves*. I can hold you by holding your hand, and I can see the coin's circularity by seeing what is visible from here (namely the coin's elliptical profile). The question is, how do we this? In this paper I address this question, and I try to show that the two-dimensionality of perceptual content is no obstacle to a

¹ One philosopher who does accept the duality of perceptual content in my sense is David Chalmers (2005). He and I agree that perception is in an important sense a "two step" phenomenon. There are important disagreements between us, though, some of which I touch on in this paper (not always drawing explicit connections to his work).

direct realism worth defending. But first, I turn (in the next three sections) to a consideration of views that deny the two-dimensionality of perceptual experience.

1. Denying perspectival content

Some philosophers reject the perspectival dimension of perceptual content. Peacocke (1983) has pursued this strategy. He starts from the thought that one experience cannot present the tomatoes as the same in size *and* as different in size, that one experience cannot present the coin as both circular *and* elliptical. He grants, though, that there is a genuine phenomenal basis to the apparently perspectival aspect of perceptual content. After all, there is some sense in which the tomatoes look different in their size (in their *extensity in the visual field*), even when we see that they are the same in size. He proposes that the elliptical perspectival shape of the coin, and the merely apparent difference in size of the two tomatoes, are not ways the experience presents things as being, but are rather qualitative or sensational properties of the experiences themselves. The experience of the same-sized tomatoes is one in which there are visual field properties that differ in respect of size. Likewise, the experience of the round coin may involve the instantiation of an elliptical visual field properties. Such sensational properties of the experience thus contribute to what it is like to have the experience but without being features of the way the experience presents the world as being.

What this view captures is the internal connection between actual size and shape and apparent size and shape. Elliptical is the way the profile of a round coin looks when seen from an angle; same-sized tomatoes take up more or less of the visual field depending on their distance from the perceiver. We are not forced to accept Peacocke's

account of this internal connection, however. Considerations of coherence would give us reason to doubt that an experience could present an object as circular and elliptical at one and the same time; but there are no similar impediments to an experience's presenting an object as circular but as looking elliptical. Why not say that perceivers are able to see both what size a thing has, and also its apparent size (how it looks with respect to size from here), or to see a thing's shape, and its apparent shape.² There's no contradiction here. Although there is, as we have seen, a looming question. We can ask: in what does the *visual* experience of the coin's circularity consist if not in the fact that it, the coin, actually *looks* circular from here? Can we maintain that we see the coin's circularity when the coin is presented by means of something elliptical?

Another philosopher who denies perspectival content is A.D. Smith (2000). He insists that coins, when seen from most angles, do not look elliptical. This is simply not true, he says. They look round, or perhaps "round and tilted away from you" (172). Smith does not deny – any more than Peacocke did – that there is a definite experiential basis to the claim that the coin presents itself by means of an elliptical perspectival property. He acknowledges, reasonably, that when a student insists that the coin *looks* elliptical to us even though we *know* it is round, she isn't simply mistaken. She is calling attention to a salient aspect of our distinctively visual experience when looking at the coin.

However, Smith does want to deny that this phenomenal aspect corresponds to the way the coin looks. He stakes out a position very much like Peacocke's. What is elliptical

² Michael Tye (1992; 2002) has developed a related criticism of "non-representationalist" accounts of perceptual experiences.

is not the way the coin looks, but a sensation, that is to say, a quality of our experience of its appearance. What justifies this claim?

Smith emphasizes that one feels no inclination – not even “the tiniest bit” -- to take a coin *to be* elliptical when one sees it (182). He asserts that “After all, if something really does perceptually look elliptical to me, I shall, if I notice the thing, and if I have no countervailing information to hand, take the thing to *be* elliptical, for I have nothing else to go by.” This may be so. However, it leaves open that what explains the fact that normal adult perceivers feel no inclination to take the coin to be elliptical on the basis of its visual appearance is the fact that they do have an abundance of countervailing information. Most of us know that coins are round, and we know that round, flat things change their perspectival shape as our spatial relation to them varies. As a matter of fact, there is considerable evidence that genuinely naïve perceivers – perceivers truly lacking the relevant countervailing information – *would* take round coins presented at an angle to be elliptical. Cheselden (discussed in von Senden) describes the astonishment of a young boy who had undergone cataract surgery at the way a coin changed its shape as it moved; Helmholtz describes a similar case of a boy astonished at the changing shape of a locket. More recently, in a similar vein, Valvo describes a patient who, shortly after undergoing surgery to remove cataracts, perceived as black holes what he later found out to be windows of a house across the street. Large objects far away looked to this post-operative patient like small holes nearby.³

³ Smith (2000, 35-47) takes pains to distinguish a merely evidentiary from a genuinely perceptual sense of “looks” or “appears”. Relying on this, it would seem, one can grant

Smith gives an additional reason for thinking that coins don't usually look elliptical (or, *mutatus mutandis*, that the tomatoes don't look different in size). The quality of our visual experience that leads theorists such as myself to say (mistakenly, on Smith's view) that the coin looks elliptical is one that, obviously, changes as our spatial relation to the coin itself changes. But normal perceivers do not experience the shapes and sizes of what they see to be changing. This shows that our experience of ellipticalness is not an experience of the coin's *being* elliptical.⁴

I agree that mature perceivers can experience a coin as elliptical without coming to think that it is elliptical. I grant, after all, that there is genuine size and shape constancy. But this does not support Smith's claim that the coin does not look elliptical. The question after all is not whether there is size and shape constancy – I grant that there is – but whether there is also *inconstancy* with regard to apparent size and apparent shape. As far as I can tell, nothing Smith says rules that out this possibility.

that the coin doesn't look to be elliptical in an evidentiary sense of "looks" while holding to the claim that it looks elliptical in a perceptual sense. Given this, it is unclear to me why Smith thinks that how one is inclined to take things to be is relevant to the question of how one experiences them.

⁴ Kelly (2005) also lays great stress on the distinction between something's looking F and looking *to be* F.

2. Denying perceptual constancy

So much for views that deny or minimize the significance of the fact that we can see the apparent difference in size of the tomatoes at different distances, or that the silver dollar looks elliptical when we see it from an angle. Another strategy for explaining away the apparent duality of content is that favored by the sense-datum theorists of old (e.g. Ayer 1955): these thinkers do not deny that the coin looks elliptical; rather they deny that it actually looks circular. They don't deny that there is an apparent difference in the size of the tomatoes, but they deny that there is any sense in which the two tomatoes really *look* to be the same size. The basic claim here is that we go beyond what is given to us in perception when we make claims about the shape and size of mind-independent objects.

Now, this line of criticism can be construed in two ways. On a first, epistemological, construal, the claim is that our experience fails to provide sufficient rational justification for the judgments we make about what is there in front of us. This is an important line of consideration, but it is not directly relevant to our discussion. Our focus here is not epistemological, but phenomenological. I take it be phenomenological bedrock that we at least experience the coin's circularity, or the sameness of size of the tomatoes, when we look at it (or them) from a given vantage point. The circularity and sameness of size at least seem to be sensibly present even if we do not take ourselves to *see* them outright. Nor is it the case that we merely *think* or *judge* or *infer* the presence of these features. There is a difference between thinking something is present and experiencing it as present. The way in which the strictly unperceived circularity of the coin is present is decidedly not the way inferred objects are judged to be present. This is demonstrated by the fact the coin would look no less circular even if one knew, for a fact,

that it was elliptical. This is a familiar occurrence when we look at pictures. What is drawn is an ellipse, but we experience it as the outline of a plate, say. We know that, in fact, there is no plate, and so nothing circular, but we still *see the elliptical form as circular*.

Our task, then, is to understand how the tomatoes can be presented in experience as visibly different in respect of size and as looking the same in respect of size. In what does your sense of the coin's circularity consist, when it looks elliptical too? How can the tomato look both circular and elliptical? How can we account for the full-blooded duality of perceptual content?

3. Merleau-Ponty's strategy

Let's consider a different strategy for denying the two-dimensionality of experiential content. Instead of denying one or the other aspects of perceptual content – the perspectival or the perspective-neutral dimensions – this strategy denies that both aspects of content can be present in consciousness at once.

Sean Kelly (2004), developing ideas of Merleau-Ponty (1945/1961), grants that we *can* see a circular coin as elliptical and that, in that sense, it can be true to say that it looks elliptical. There is, as he says, a perceptual attitude we can adopt in which the coin really does look elliptical. However, Kelly stresses, this is not our normal attitude. In the normal, engaged attitude, we experience the coin as circular, and it is not the case that the coin looks elliptical when we experience it as circular. The engaged attitude is primary, for Kelly. To see the circular coin as elliptical we must disengage or detach ourselves from the world, or from our engagement with the world. It is quite a feat to be able to do

this at all. Indeed, Kelly entertains the possibility that we are able to do this only thanks to the invention of the technique of artificial perspective in the Renaissance. Whether or not this is true (and I find it impossible to believe that it is), the main point, for Kelly, is that we can't occupy detached and engaged attitudes at once, so we can't experience the coin, at once, as elliptical and circular. In this way Kelly, and Merleau-Ponty, discharge the appearance of duality of perceptual content.

How, according to Kelly and Merleau-Ponty, do we experience a coin to be circular when our visual access to it is confined to a limited perspective, a perspective from which the visible profile of the coin is in fact elliptical? Merleau-Ponty's account, as interpreted by Kelly, goes like this: when you see a coin from an angle, and you are in the engaged attitude, what it is in virtue of which the coin looks circular to you is that you implicitly register the fact that your vantage point is *non-optimal*. Your visual sense of the coin's circularity consists in your appreciation that you would need to move, in such and such ways, to get a better view of the coin's shape. In this way, Kelly explains, the view allows context to play a *normative* role in perception; we experience the coin in a context, and we have a kind of practical understanding that the context guides us in certain ways.

This view faces a serious problem. We can agree that there may be more or less optimal vantage points on a coin, but only relative to certain purposes. For example, if I want to read the coin's face, or pick it up, or determine its condition, or the year it was minted, then there are more or less optimal vantages points. But vision itself is not relative to certain purposes; seeing is all-purpose. For this reason it isn't clear why we should think that some perspectives on the coin are, as it were, visually privileged. The

shape of the coin fixes ways it would look from an infinite number of perspectives, and in so far as we do experience the coin's circularity and flatness, then we experience it as having a shape such that its appearance would vary in an infinite number of ways as we move in respect of it.

Putting this serious problem to one side, Kelly's Merleau-Pontyan *contextualism* would seem to be vulnerable to another criticism: it supposes that we are *aware*, when we are in the engaged attitude, not only of the coin and its qualities, but of the context and the extent to which our relation to what we're looking at is more or less optimal. But surely if we balk at the idea that we are aware of the coin's perspectival shape when we see it from here, we should also balk at the thought that *we are aware of* the ways our movements would improve our view of the coin.

If one does accept that the engaged attitude makes room for this sensitivity to variations in context, then what principled reason is there to hold back from saying that perceptual engagement with the world makes room not only for attention to how things are, but also to how they look from here? Indeed, I wonder how the proponent of contextualism can *avoid* saying this. For, we must ask, what is it about my context vis-à-vis the coin that enables me to tell my viewing angle is non-optimal? Surely it is just the way the coin looks, e.g. that it looks *elliptical* from here; it is this fact that would guide me to move in order to get a better view of the coin's circularity. What other contextual fact is available to play this "normative" role?

The contextualism we are considering seems to collapse into the view I myself want to advocate. On this view, the way the coin's circularity is present to me when I view it from a single point of view, is *as available to me* thanks to my implicit

understanding that were I to move in these and these ways its apparent shape would change in correspondingly. I *encounter* its roundness *in* encountering its elliptical apparent shape together with a practical understanding that that shape depends on my spatial relation to the coin and would, therefore, be modified by movements. From this standpoint I can comfortably admit that, most of the time, I'll attend to the coin and not pay attention to how it looks from here. And I can admit that if I were to attend to how it looks from here, then I would need to divert my attention from how it is in itself. But these admissions do not require us to suppose that, in thus shifting our attention back and forth, we move between incommensurable *attitudes*. Experience contains within it precisely two aspects, or dimensions, to which we can turn our attention.

This brings me to a third criticism of the Merleau-Pontyan contextualism. Perhaps what underwrites the intuition that the engaged and detached attitudes exclude each other is an idea (entirely correct in itself) about limits of attention and the background. When I focus my attention on the coin's shape, the elliptical appearance property that would be available to me as the focus of my attention were I to shift to the detached attitude, is present in experience only as a feature of the background. If this is right, then the opposition between engaged and detached attitudes is just a symptom of the figure/ground structure of conscious attention. Understood this way, however, we have to recognize that the contextualist's worry isn't so much directed to anything problematic about painterly detachment and the engaged attitude. We can grant, then, that you can't attend to shape and perspectival shape, or color and apparent color, *at once*. What explains this is not that each of these qualities is proper to what are in effect incommensurable *attitudes*, but rather that there are basic limitations of attention. True,

you can't attend to the shape and apparent shape of the coin *at once*, but you can't attend to its shape and size at once either, or its shape and color.⁵ Importantly, in so far as the world is available to me now in my visual experience, it is available to me both as it is in itself apart from my perspective, and as reflecting my perspective. Perceptual experience retains these two dimensions of content.

4. Perceptual constancy and amodal perception

We have been considering examples of what psychologists call *perceptual constancy*, specifically, size and shape constancy. Traditionally, psychologists characterize perceptual constancy this way: we experience objects as unchanged or as constant and stable in their perceptual properties despite the changing character of perceptual sensation or perceptual stimulation. When you approach a tomato, it looms in your visual field – as you get nearer, the size of the image of the tomato on your retina increases; it doubles in size as you halve your distance – yet you do not experience the tomato as changing in size. As you change the angle at which you view the silver dollar, the shape of the image of the coin on your retina changes, yet you do not experience the coin as changing in shape.⁶

⁵ This is an important lesson of empirical work on so-called change blindness. The color of an object can change without your noticing it, even if you are intent on focusing on the object. See O'Regan and Noë 2001 for discussion.

⁶ This is roughly how the philosopher Smith characterizes perceptual constancy (Smith 2000, 170).

We are now in a position to notice that this way of describing the phenomena of perceptual constancy is unacceptable. The question is not: how can our experience continue to present things as unchanged when the character of our experience is continuously changing. Rather, the question is or ought to be: how can our experience present things as unchanged when it manifestly presents them as changing. Perceptual constancy has the character of a paradox: you experience the size as unchanged despite changes in the apparent size; you experience the shape as unchanged despite the fact that the coin looks different in respect of shape!

Let's introduce a new example of perceptual constancy: color constancy. This example will help us to better appreciate the dual aspects of perceptual content and so the apparently paradoxical character of perceptual constancy.

When you leave your study, illuminated with fluorescent light, and head out into the day, you don't notice that the book you are carrying, or the skin tone of the person you are talking to, has changed. In a sense, of course, they have. It is this sort of change that underwrites our judgment that, for example, no one looks good by the light of the New York Subway, or the color of the pigment you would use to match the book outdoors is different from the color of the pigment you would use to match the book indoors. Of course, we may not notice these changes in color. In part this is because our

very grip on color is a grip on something that changes dynamically in determinate ways, as we move, or as lighting changes, or as the colors of surrounding objects change.⁷

Reference to painting here doesn't mean we are taking up the painter's detached attitude. Whenever color judgments matter, it is important for us to pay close attention to the way colors change as lighting changes. If you are buying paint for your house, you'd be advised, for that reason, to bring your paint chips home and look at them in the prevailing lighting conditions at home. To do this would not be to lapse into the detached attitude of the Renaissance. It would be to take up the concerned attitude of a home owner, an attitude that, presumably, we share with our prehistoric forebears.

We are able to experience a surface as uniform or regular or stable in color even when it is visibly variable and differentiated in its color. Consider a new clean car on a rainy afternoon (as discussed in Broackes 1992). The car reflects lights and surfaces, and yet you can experience it as uniformly colored even as it thus reflects highlights, etc. Or consider a wall that is only partly cast in sunlight. Despite its looking brighter where the sunlight hits it, you needn't experience the color there as different from that of the rest of the wall (even though differences in brightness make for differences in color).

Perceptual constancy phenomena such as these exemplify a broader category of perceptual phenomena, one that philosophers have tended to neglect: amodal perception. Look at an apple on the table in front of you. You have an experience of a voluminous

⁷ See Noë 2004, chapter 4, for extended discussion of this. I there argue that it is a myth that colors present themselves to us in experience as *simple* and *manifest*. Colors have hidden sides to them, just as shapes do.

whole. In fact, you can only see the facing surface of the apple. It doesn't seem to you that you actually *see* the whole of it. After all, what could be more salient than the fact that you only see its visible features, its facing side. And yet, in seeing what is visible, you have a sense – a visual sense – of the presence of the apple as a whole.

Similar considerations go for a cat, whose body is hidden by an obstacle, but whose head and tail show. You can *visually experience* the head and tail as belonging to the same cat, even though you can't actually see the part of the cat to which they are joined. Here it doesn't seem to you, erroneously, as if you see something you don't. Rather, it seems to you as if you are aware, in a perceptual modality, of something that is plainly out of view.⁸

We can give substance to the sense that your awareness of these items is not only perceptual, but visual. The distinction between a thought-state and a perceptual one is clear enough. As we have noticed, you don't merely *judge* the apple to be present; it looks to you as if it is. This difference is grounded in the way in which perceptual experiences, but not non-perceptual ones, are highly responsive to movements of the body and changes in the environment. In perception, your relation to the perceived features is *sensorimotor*. What makes the perceptual relation visual, as opposed to some other perceptual modality, is the precise character of the mediating sensorimotor relation.

⁸ For this reason it is misleading to name this phenomenon *amodal* perception. This term marks the fact that the experience we have of the apple's backside is an experience of what we plainly perceive in *no* modality. Nevertheless, the *way* in which the apple's backside is present (but unseen) is clearly visual.

In the visual case, movements of the eye and head play a special role in modulating sensory stimulation. In so far as you now stand in a relation to the apple and the cat that is mediated by this sort of pattern of dependence of your visual experience on what you do, and in so far as you implicitly understand this to be the case, then to that extent, it seems, your relation to these items is visual. As has been argued elsewhere (O'Regan and Noë 2001; Noë 2002, 2004, 2005), it is the obtaining of this kind of *sensorimotor* relation, and *the perceiver's implicit understanding that these relations obtain*, that is the ground of the claim that, in seeing the visible side of the apple, or the visible parts of the cat, you have a *visual* sense of the presence of the strictly unseen, or strictly *invisible* parts of the cat.

To summarize: in what does your perceptual sense of the presence of unseen items consist? It consists in your practical knowledge of how to bring those unperceived items into view by movements of the body; in your skill-based sense of their availability.

Now, I propose that perceptual constancy phenomena are themselves examples of amodal perception. The relation of the coin's circularity to what you see, or of the tomatoes' sameness of size to how things look to you, is the same as that of the relation of the apple's voluminousness to what you see. The size of the tomatoes, the actual shape of the coin, the actual color of the wall, these are present in your experience of the tomatoes and the coin and the wall the way the unseen parts of the cat or the voluminousness of the apple are present in your experience. They are present *as absent*, *but as available to perception through appropriate movement*.

5. Enacting direct perception

Presence in absence is not illusory presence; it is rather a special kind of availability. The world is present, in perception, not by being present (e.g. represented or depicted) in consciousness all at once, as it were, but by being available all at once to the skillful perceiver. And different items are available *in different ways*, depending on the kind of bodily, sensorimotor relation that we hold to them. (For more on this, see Noë 2004 and Noë 2005.)

Perception is the encounter with the world from a point of view. The limitations of perspective are therefore constitutive of the perceptual relation. There is no seeing a peach from all sides at once; there is no seeing a wall as it would look in every illumination; there is no seeing the coin's circularity from an angle. The availability of the coin's shape, or the peach's body, or the wall's color, is made possible thanks to our implicit understanding of the way the appearance of these objects would change as we move or would move in relation to them.

The sense-datum theorist was right to this extent: perceptual access to the world is mediated by how things sensibly appear. But this is compatible with direct perception. For what mediates our perceptual relation to the world is only our exploration of the world. The status of the claim that perception is mediated is exemplified by the fact that you can't see the peach from all sides at once. I can hold you, but I can't hold you by holding every single part of you. I can see you, by seeing your surface. You and the world can be available to me thanks to my sensorimotor understanding of the way my contact with you is a kind of contact with the world that is *beyond view*.

Crucially, what mediates the perceiver's relation to the apple, the peach, the wall, the coin, the color, is not a *sensation* – if by this we mean mere qualities in consciousness or qualia – but rather skillful contact with the perceptual world. What explains the elliptical appearance of the coin are not facts about my qualia, but rather facts about how the coin looks from here, facts that are determined by the coin's shape and my spatial relation to it, by the coin's place in its environment. What mediates our relation to the world in perception, then, is the world, and what we do or can do. So the two steps can, in a way, collapse into one. Perceiving is *exploring the world*. It is a temporally extended activity. What we call *seeing the apple* just is an episode of exploration. And so we can say that we *enact* the perceptual world by skillful exploration. On this way of thinking about perceptual experience, perceiving is not a way of representing, it is a way of gathering or assembling content.

6. Sensation in perception

From this standpoint, we can venture to criticize A.D. Smith's strategy for resisting the Argument from Illusion, laid out in his recent book (2000). The task for direct realism, according to Smith, is to show that perceptual awareness of objects is not mediated by prior awareness of mere sensation (or sensory qualities). Smith argues that the phenomenon of perceptual constancy can enable one to do this. For when "perceptual constancy is in operation, although perceptual sensation is changing, what I am aware of, in the most basic and immediate sense, does not appear to change. (180)" When you approach an object, your perceptual sensations vary, but your experience of the size and

shape of the object do not vary. It follows, then, that what you take yourself to be aware of, is not the sensations themselves.

This argument rests, as I hope will be clear, on a mischaracterization of perceptual constancy. I have argued that when you are perceptually aware of an object, you are, *inter alia*, aware of the way it looks. The way it looks, however, *does* vary as you move in relation to it. For this reason, perceptual constancy cannot play the role Smith wants it to play, that, namely, of distinguishing object-directed perceptual states from merely qualitative (sensational) states. It is not that perceptual constancy somehow fails to enable us to differentiate between the object of awareness, with its stable perceptible properties, and the variations in the way the object looks from here or there. The problem, for Smith's account, is that the variations in how the object looks are not a matter of "mere sensation." There does not seem to be a place, in the analysis of perception, for what Smith calls perceptual sensation.

Why does he think that there is? Enter the Argument from Illusion. According to Smith, the Argument from Illusion demonstrates conclusively ("beyond any shadow of a doubt") that the very same perceptual sensations are instantiated in veridical and nonveridical perceptual experiences and they are instantiated in the very same way (Smith 2000, 65). The Argument is as follows: Consider the visual experience of a white wall looking yellow. In so far as the wall really looks yellow, then, in seeing the wall, you experience yellowness. But this yellowness that you encounter is just the same yellowness that you would have encountered were the wall really yellow. What else could explain the fact that the wall looks to be, precisely, *yellow!* But the yellowness you

experience looking at the white wall, is not a property of the wall. The wall is white and the yellowness is a property of your experience.

A philosopher might object that from the fact that you misperceive a white wall as yellow, it does not follow that anything really is yellow. This “sense datum inference,” as Smith calls it, is fallacious (Austin 1962; Putnam 1999). Smith insists that such an objection cannot be taken seriously. “When, in the situation in question, I describe the wall as yellow, I do not just pick a colour term at random. Surely yellowness is appropriate to giving expression to my perceptual state only because that is the colour *I am aware of*. I say “yellow” because I see yellow. When something appears yellow to me, it is, or could be, with me visually just as it is when I veridically see something that really is yellow. The same “sensible quality” is present to consciousness in the two cases. For were it not, why should I make reference to the same colour in the two situations? (37)” Smith is insistent that we need “to recognize certain experiential facts,” to wit: “When a wall perceptually looks yellow to you, it may be with you experientially just as when you veridically see a yellow wall. The experiences are, or may be, qualitatively identical. The Argument [from Illusion] claims that the only way to do justice to this fact is to recognize that a veridical and a matching a illusory experience have a *shared sensory character*. When a wall perceptually looks yellow to you, a certain sensory quality is realized in your experience whether or not the wall is yellow.” (40) According to Smith, we can be as certain of this as we can be of anything.

Let us agree with Smith that there is such a thing as a white wall’s *actually looking yellow*. To explain this phenomenon, must we accept Smith’s thesis about sensations? Is it true, as Smith asserts, that the *only way* to do justice to the possibility of

illusion is by supposing that the same perceptual sensation is realized in an illusory experience and its corresponding veridical experience?

In fact, we have already considered reasons to be skeptical of this qualia thesis. Generally speaking, how things look is not a matter of “mere sensation” or qualia. How things look is a feature of the way they are. There is no more reason to think that a wall’s looking yellow is a fact about my *perceptual sensations* or qualia than there is to think that a coin’s looking elliptical from here is a fact about my sensations. To say this is not to deny “the experiential facts.” It is to deny that mere qualia play a role in fixing or constituting those facts. Let us ask: What is it in virtue of which the round coin looks elliptical? Answer: geometrical facts about my spatial relation to it, not facts about my internal, nonintentional qualitative states.⁹ An exactly analogous reply is ready to hand in the color case. What is it in virtue of which the white wall looks yellow? Answer: facts about the character of the illumination, or about the chromatic properties of nearby surfaces (e.g. the wall might be reflecting the yellow from the mustard flowers outside the window). Someone might object that one would have the experience of yellowness so long as one’s brain was in a certain state regardless of any environmental conditions. But what does this show? Not that we must appeal to sensations or qualia to explain perceptual experiences, but that we must appeal to brain states.

⁹ My claim is not that geometry is sufficient for experience alone. A blind person wouldn’t experience the coin’s shape whatever spatial relation obtains between him or her and the coin!

I want to argue, further, that sensations or qualia, in Smith's sense of *perceptual sensations with no intrinsic intentionality*, can play *no role* in explaining the occurrence of this sort of nonveridical experience. This is so because sensations, thus construed, are neither necessary nor sufficient for the experience of the presence of a perceptual quality. That they are not *necessary* is illustrated by such simple facts as that blinking, which interrupts sensation, does not interrupt perceptual *experience*. A more robust example is provided by amodal perception. You have a sense of the presence of the apple as a whole, even though you only see its facing side, that is, even though only the facing side gives rise in you to *sensation*. That perceptual sensations (i.e. mere sensations with no intrinsic intentionality) are not *sufficient* for the sense of presence of a feature, is illustrated by a consideration of someone who has adapted to reversing goggles. The goggles affect incoming stimulation so that an object on the left stimulates the eye and brain as an object on the right normally would. That is, the object gives rise to a *right-sensation*, even though it is on the left. After adaptation, however, the object looks as though it *is* on the left, even though it continues to produce right-sensations (that is, sensations that would normally be produced by an object on the right). The quality experienced varies even though the sensation does not.

If the occurrence of mere sensation does not explain the experience of location of an object in the visual field, what does? An object on the left is one whose appearance changes in precise ways as one moves in relation to it. What fixes the experience of an object *as on the left* is not the intrinsic character of sensations, but the perceiver's implicit grasp on the way *how the object looks* would vary as he or she moves. Adaptation to reversing goggles, then, is a matter of figuring out the patterns of sensorimotor

dependence governing your relation to the object, not a matter of having certain kinds of feelings (O'Regan and Noë 2001; Hurley and Noë 2003; Noë 2004).

And so for other kinds of perceptual content. In what does your visual sense of the cubicalness of a cube consist? It consists in your implicit expectation that its appearance would change in the characteristically cubical way as you move in relation to it. Sides come in and out of view, vertices come in and out of view in the ways definitive of cubicalness. Importantly, there is no *cube* sensation, just as there is no sensation as of leftness. To experience something as a cube, or as on the left, is to experience it as exhibiting a characteristic *sensorimotor* profile. Likewise, the hidden parts of the apple and the cat are present, in vision, as accessible, not as *given* (as it were).

What these considerations show is that how things visually seem to be is not explained by appeal to sensations (at least if we understand these nonintentionally as Smith does).¹⁰ We can't explain the qualitative features of experiential states (veridical or nonveridical) by supposing that they depend on the instantiation of sensations in consciousness.

If it is possible to do justice to the possibility of illusion without appealing in the way Smith does to mere sensation, then the mere possibility of illusion doesn't imply anything about the role of qualia or perceptual sensation in illusion. The Argument from

¹⁰ If we allow for a conception of sensations as themselves intentional, that is, as world-directed, then it is unclear what explanatory work they could be called on to do. For sensation thus conceived would in effect be perceptual experience, but that is precisely that for which we seek an explanation.

Illusion, it would seem, leaves us where it starts, with the observation that illusion is possible, i.e. there needn't be anything in the character of one's experience to enable one to tell whether the qualities one seems to see are really present or not.

7. Real presence

It will be objected that I haven't done justice to the main challenge posed by the Argument from Illusion: It is possible to experience the presence of yellowness when yellowness isn't physically present. We can make sense of the possibility of a white wall's looking yellow for no other reason than that the brain is in the "yellow state" even though lighting is normal (etc). In such a case, then, the yellowness belongs to one's state of consciousness, not to the perceptual world. What we need, Smith, for example, has insisted, is an account of this experiential fact; what we need is some account of how this fact can be reconciled with direct realism.

Let us turn back to the experience of yellow when nothing is yellow. This is, in fact, a special case of the more general phenomenon of experiencing what is absent as if it were present. Presence in absence, or amodal perception, is, as we noticed earlier, a hallmark of *normal, veridical* perception. When you look at the apple, you have a sense of its presence as a voluminous whole, even though you only actually see its facing side. When you look around your room, objects partially block each other from view, but you have a sense of the presence to vision of those partially blocked objects as wholes. It isn't as though you take yourself to see what is out of view; rather, you take yourself to experience as present what is, evidently, out of view. Phenomenologically, as we have considered, this sense of the presence of what is out of view is an experience of it as

accessible. The apple is accessible, and not just in any old way, but in the distinctively bodily way characteristic of perceptual consciousness: movements of the body bring unseen things into view; your relation to the unseen is mediated by sensorimotor relations with which you are familiar.

Can we extend this sort of account of normal, veridical presence in absence to hallucinatory experiences of presence such as that of the yellowness of the white wall? You might think that we cannot do this, and for the simple reason that the yellowness of the white wall is not experienced amodally as present as absent, but rather as *modally* present, that is to say, as in full view. In the visual experience of a white wall as yellow, one doesn't experience the wall's yellowness as merely available, but rather, as witnessed. The contrast here, then, is between the *modal* nonveridicality of the hallucinatory experience of the white wall as yellow and the *amodal* nonveridicality of the experience of what is a façade as the experience of a whole building. If I mistake a stage building for a real building it will not be because I actually take myself to *see* (say) that there is a saloon beyond those swinging doors, or a bedroom behind that curtain. Nor do I merely *think* that they are there. Rather, I visually experience them as present, although out of view, by drawing on a repertoire of sensorimotor anticipations. In the case of the yellow seeming white wall, though, I *do* take myself to *see the yellowness*; it is not that I take myself merely to have access to it by movements of the eyes or head, by the exercise of sensorimotor anticipations.

It is certainly true that the account of perceptual presence I have laid out here (and elsewhere, e.g. Noë 2004, 2005) seems to trade on a contrast between what is present, *simpliciter*, and what is present, but absent, out of view, merely available. There are,

however, good reasons to doubt that this distinction is hard and fast. Once we move beyond this contrast, we will be able to see how we can extend the enactive account of perceptual presence to the hallucinatory case of the white wall as yellow.

At first blush we think that the face of the apple is present, *simpliciter*, whereas the sides of the apple are only visually present *as absent*. But consider: as a matter of fact, you don't visually experience the *whole* of the facing side of an apple, when you experience the apple (when you see the facing side and experience the visual presence of the parts that are out of view). After all, you cannot simultaneously attend to every spot on the apple's facing surface. You do have a sense of the presence of the apple's facing side *as a whole*, just you have right now a sense of the presence of the apple as a whole. But not because you actually *see* the whole of the facing side, any more than you actually *see* the whole apple. It turns out – or so I want to argue – that there is no place to draw a sharp line such that what is on this side of the line I really experience and what is on that side of the line I experience as present without really seeing it.

The facing side of the apple, no less than the hidden parts, are present to you, now, only as available. Movements of your eyes and your attention can bring parts of the facing side into view, just as movements of your eyes (etc) can bring parts of the hidden sides into view. Of course there are differences in your visual relation to the visible and hidden parts of the apple. But these differences are matters of degree, not kind. Your relation to the apple is mediated by patterns of sensorimotor dependence such that movements of your eyes and movements of the apple produce sensory change.

I mentioned above (in section 5) that although there is truth in the two-step approach to perception – perception is a process of learning how things are from how

they perceptually seem – there is a sense in which the steps collapse into each other. We are now in a better position to understand this collapse. Exploring how things perceptually seem to be *just is* exploring how they are. For a skillful perceiver – that is, for one who understands the way sensory stimulation, how things look, changes as you move, or as objects move – the encounter with appearance is an encounter with how things are independent of mere appearance. *All perceptual content is present in experience as available to us thanks to our ability to probe.* The distinction between the given and the merely available is erased and we are left with the idea that perceptual awareness of an object with its stable properties is, in fact, a dynamic activity of skillful interaction with things around us. Perceiving isn't representing, or even presenting; it is *enacting* perceptual content — that is to say, making contact with the world through skillful exercise.

As with objects, so with colors. To the facets of an object there corresponds, in the domain of color, the variations in the way a color looks as illumination changes or as colors of surrounding objects change. Just as there is no *sensation* that corresponds to the perceptual awareness of the object as a whole, so there is no *sensation* that corresponds to the perceptual awareness of a color as a whole, however much this goes against the grain. Yellowness is multi-faceted and multi-dimensional and yellowness is never present, in experience, all at once, for the yellowness of a surface is only revealed in the way the surface changes its appearance as color-critical conditions (e.g. lighting) change (Noë 2004, esp. chapter 4). Yellowness is no more visible than is the panoramic object entirely visible (nor is it less visible).

One more example: you are in a stadium and you look out across at the facing stands. You have a sense of the perceptual presence of thousands of faces, even though, of course, it doesn't actually seem to you as if you see them all in sharp focus. You experience the scene in all its detail, to be sure, but not in the sense that you *represent the detail in consciousness*. It certainly does not seem to you as if you do that, and empirical psychology demonstrates decidedly that you do not.¹¹ In what sense are you visually aware of the detail if you do not literally see the detail? The crowd is present in that it is available to you, and it is available to you in distinctively *visual* ways; it is accessible to your directed looking (the visual equivalent of touching). Colors and objects are like crowds: they transcend what is given.

The Argument from Illusion does not show that veridical and nonveridical experiential counterparts have the same (nonintentional) sensational content. For perceptual phenomenology does not supervene on mere (that is, nonintentional) sensational content. Nor does it show that we are *aware of the same thing* in veridical and nonveridical cases. In the perceptual case we are aware of yellow; in the nonveridical case, it is merely *as if we are*.

This *as if* similarity needs explaining, and we can explain it: we spontaneously draw on the same sensorimotor expectations in both cases. Experience is a temporally extended activity of skillful-engagement, or apparent engagement, with the world. It isn't that perceiving yellow and having the corresponding hallucination of yellow is a matter of being in the same state; it is a matter of *acting, or being disposed to act, as if we were*.

¹¹ Again, I am thinking of work on change blindness.

This point here is close to that of Martin (2004), who claims that all that we can say about an hallucination is that it is, for the subject, indistinguishable from the corresponding perception. I'll return to Martin's reasons for claiming this. For now, I want to stress that I try to go farther than Martin is willing to go in suggesting that, without supposing that there is common content, there is something more to be said about *why* the experiences are indistinguishable. The hallucinator acts out the same *sensorimotor repertoire* as the perceiver.

The reader may worry that although this way forward manages to block the Argument, it does so at too high a price, in a way that forfeits any claim to direct realism, that is, to the idea that experience itself *involves* situations and things. On the enactive view, after all, even in veridical perception, we do not quite see what we think we see. Isn't real presence, on this view, an illusion? The world is never really present *in our consciousness*; it is only *as if it is*. But this line of criticism trades on a misunderstanding. Real presence is no illusion. For it does not even seem to us as if the world is present *in our consciousness*. It does not seem as if we represent it all in mind. This is a theoretically motivated misdescription of what our experience is like.¹² When we perceive the apple, we do experience it as present, and it is! But this is compatible with the fact that we don't (and certainly don't take ourselves to) see the hidden parts of the apple; it's not as though we have *all of it* in consciousness (just as it isn't as if we see it

¹² I here suppose that there is a theoretically neutral way of characterizing what our experience is like. This a point with which Dennett (2001, 2002) would strenuously disagree.

from all sides at once). The world is *really* present in veridical experience, but not as *represented by us*, not as *in our heads*; rather, the world is present as *out there* and as *available* to our inspection. In nonveridical experience, the world isn't really present in this way, although it seems that it is.

8. The possibility of hallucination

In a recent discussion of perceptual hallucination, Michael Martin writes: “At present, I have good reason to suppose that I am seeing a London street pretty much as it is...Nonetheless, as far as I can tell, it seems a genuine possibility that I could have been in a situation which was not one of actually perceiving my environment for how it was but which I would not have been able to tell apart from this, my actual situation, just through introspection and reflection on my experience. Such a case would surely be a perfect hallucination of the kind of scene that I am perceiving, as things stand, for what it is” (Martin 2004).

If one thinks of perceptual experiences as self-standing interior states, like pictures in the head, or inner projections, say, then it does seem like a “genuine possibility” that one could be in a situation which was not one of perceiving things around one, even though it was for one just as if one was. Given such a conception of experience, the burden would seem to be firmly on the direct realist to establish that experience, so conceived, could be a form of contact with the world around one. Much philosophical research in this area is spent trying to shift the burden of proof from the shoulders of the direct realist to one who denies that experience can be a kind of direct contact with the world (e.g. Martin 2004).

There is another way of thinking about perceptual experiences. On this alternative conception – which is no less naturalistic in its starting assumptions – the paradigm of perceptual experience is active touching. Consider a blind person’s experience of the texture of the sidewalk as he or she tap-taps along with a cane. He or she experiences the sidewalk and its texture, but not by forming a picture of it, in the head, as it were. The experience, rather, is the temporally extended activity of probing, one that is constituted not only by what the perceiver does, but by the way the world resists or even acts back. On the more traditional Cartesian conception, experiences are self-standing and internal; they are states brought about in us by stimulation from the environment. On the alternative way of thinking of perceptual experience, experience is not something that takes place within us, but something we do, in the world.

From the standpoint of the alternative conception of experience, which has been taking shape in this discussion, the difference between seeing a yellow wall and experiencing a white wall as yellow (in the absence of the lighting conditions or whatever that would explain the yellow appearance) is comparable to the difference between, say, dancing to music and dancing without the music’s guiding influence, or between moving your hand along a surface, in a way that is guided by the contours of the surface, and moving your hand in the same way without the surface’s guidance. This comparison points to what a perception and a corresponding hallucinatory experience would have in common: the exercise of the same bodily, sensorimotor skills. And to this extent, the comparison directs us to what would explain their indistinguishability. It also reveals why perfect hallucination is a *near impossibility*. It is practically impossible to exercise dynamic sensorimotor interaction with the world in the world’s absence; the chances of

perfect hallucination are like the chances of moving one's hands as if one were guided by a surface when one is not; they are almost nil.

The philosopher – direct realist and skeptic alike – will insist that what is in play, in the Argument from Illusion, is *the bare a priori possibility* of perfect hallucination (as Martin 2004 calls it). The fact that perfect hallucination is unlikely is neither here nor there. I grant this. *Perfect hallucination is logically possible*. But in granting this there is no need to make any concession to the internalist conception of experience. The conception of experience as self-standing and internal is not something the Argument helps to legitimize, but at best, something it presupposes. Without that presupposition, however, the Argument loses its point, at least to some extent. For the bare possibility of perfect hallucination had seemed to flow naturally from the starting-point conception of experience: how things are with one experientially is, on the conception in question, a matter of how things are *with one*, or rather, *in one*. The possibility of perfect hallucination had seemed, thus, to be part and parcel of our subjective, experiential lives. If we do without the Cartesian presupposition, however, then all that the bare possibility of perfect hallucination comes down to is the fact that there is nothing incoherent in supposing that things could be so gerrymandered as to give us the neural stimulation that we would receive were we dynamically acting on and being affected by the environment. The stage setting required to perpetrate perfect hallucination on unsuspecting perceivers would be quite a feat indeed! But if this is what the Argument comes down to, then it does not show that we are confined to the contents of our own subject states of consciousness; rather, it demonstrates only that the world could, possibly, conspire

to mislead us radically. The Argument thus deflated shows no more than that hallucination is possible. This is, of course, something we knew at the outset.

Perhaps it will be objected that I am ignoring the fact that we have abundant experience, in our real lives, of perfect or near-perfect hallucination. The fact that we dream suffices to show that experience is entirely a matter of what is going on within us. But this is simply not true. It is rare that dreams are truly life-like, and when we find ourselves convinced by them, as we certainly sometimes do, it is rare that their life-likeness results from their *perceptual* quality and not, say, their emotional content. Perceptual content of dreams tends to be noticeably sparse and unstable.¹³

9. Facing up to the challenge

Perceivings and their hallucinatory counterparts have in common, I have argued, not shared content, but the appearance of shared content, for they have in common the exercise of the same bodily, sensorimotor skills. It is this that explains how it could happen that the experience of nothing is for all intents and purposes like the experience of something.

¹³ Afterimages might also be a good candidate for perfect hallucination of color. But there are usually visually salient clues as to the nonveridical status of afterimages. For example, they move with the eyes, change size depending on the background, and grow more intense when the eyes are closed. Nor does virtual reality fit the bill. Virtual reality experiences (in, e.g. a fight simulator) are best thought of as veridical experiences of virtual worlds, i.e. specially engineered environments.

You may worry that this proposal is counterintuitive. How, you ask, is the “exercise of sensorimotor skills” supposed to explain why hallucinatorily seeming to see something can be just like actually seeing it? I have tried, here and elsewhere, to show how practical mastery of sensorimotor skills serves as a phenomenologically plausible determiner of the quality of experience. Indeed, the enactive conception of experience outlined here is phenomenologically plausible. It does not seem to us, when we experience the world around as rich in detail, as if all that detail is *in our minds*. The detail seems to us to be out there, available to our inspection, by the exercise of our sensorimotor skills (e.g. by moving our eyes and head).

In any case, it is no test of my proposal – which is in essence an empirical proposal – that it meet some standard of intuitiveness. For we are now in a domain where intuition has no standing. The enactive, sensorimotor proposal put forth here is surely not *more* counterintuitive than the far-fetched possibility (however “genuine”) of perfect hallucination itself. To imagine the sustained occurrence of the neural correlates of a perceptual experience over a period of time in the absence of the environmental partners to its sustained occurrence is like imagining the existence of a pattern of shadows of the ground in the absence of the material objects that cast them. We *can* imagine it. But so what?

But the defender of direct realism will not be satisfied. If the exercise of sensorimotor understanding fixes the character of experience *independently* of whether or not the object we seem to see is present, then, the object cannot after all be itself an element in the experience. As Martin has convincingly shown (2004; personal communication), in attempting to say something *positive* about the hallucinatory

experience – that is, in going beyond the merely negative epistemological fact that the experience is one that cannot be told apart from a genuinely *perceptual* one – we run the risk of demonstrating that there could be an experience *of the very same kind* as the perceptual one but that is *non-perceptual*, that is, *non-object-involving*. But that would be to make the object we seem to experience inessential to the experience itself.

One strategy, suggested recently by Putnam (personal communication), would be to insist that we can explain why it is not possible to tell hallucinations apart from veridical perceptions even though these are very different kinds of states (the latter but not the former being episodes of contact with situations and things). What explains the indistinguishability of these episodes despite their basic difference is that fact that each is, as I have suggested, an instance of the deployment of the very same repertoire of *skills of understanding* (as I will call the sensorimotor, and perhaps also conceptual, skills involved in perceiving). Now we come to Martin's worry that we have "screened off" the world itself because we have explained, or try to explain, how you can have a state like that of seeing when one is not seeing (when one fails to stand in a suitable relation to an object). But this is too fast. Putnam insists, at this juncture, that when skills of understanding are deployed in the hallucinatory episode, their deployment counts as a kind of *misfiring*. The possibility of this sort of misfiring, however, doesn't show that the larger skillful practice of exercising the relevant skills is not world-involving. Crucially, no independent reason has been given to challenge the background fact that our skillful practices are skills at home in the setting of our ongoing *exchanges* with the environment.

But Putnam's argument, however persuasive, doesn't take us as far along as we need to go. It provides a kind of transcendental argument to the effect that hallucinatory

experience is only possible against the background of nonhallucinatory experience, an argument which, if sound, provides a strong defence against a global skepticism. But what we need to establish, if we are to defend the idea that perceiving is the achievement of contact by perceivers with situations and things, is not merely that there is sometimes veridical experience, but further, that when we visually experience things, the things we experience are ineliminable elements in our experience itself. The world we see is not merely a condition on the possibility of veridical experience. Situations and things are sometimes constituents of our experiences themselves. — Can we provide for this?

To do this we might claim, as I have hinted, that the relevant skills of understanding are such that we can only deploy them in the presence of the objects to which they apply. Indeed, this is just what I want to claim. But we immediately face a problem. If this is the case, then it can no longer be seriously maintained that what accounts for the indiscriminability of my experience and the corresponding hallucination is that in each I exercise the very same skills of understanding. For, in so far as we think of those skills as world-involving, we do *not* exercise the very same skills in seeing and in merely seeming to see. Once we concede this and once we allow further that it is not the application of the skills themselves, but their *seeming* application, that does the explanatory work, then we are back where Martin predicted we would find ourselves. For now we have a putative explanation of why experience has whatever character it has that is indifferent to the presence or absence of the object with which the experience is meant to be an episode of entanglement.

Must we then settle for Martin's negative epistemological fact? If we think of experience as an *involvement* with situations and things, must we be silent about what

makes it the case that a hallucination may be for us just like an episode of seeing? I think not. Consider a point to which I will return in the next section: I have proposed that we can account for the character of our seeings in terms of ways our sensorimotor skills enable us to secure contact with situations and things. We then ask, what can explain the possibility of hallucination? If perceptual consciousness is a mode of contact with things, how can it seem to us as if we are in contact with situations and things, in the normal perceptual way, when we are not? What I have proposed is this: when we hallucinate, the very same skills are *triggered* or *called into play*; the idea is that it is this fact that explains why we think of ourselves as in touch with things when we are not in touch with them. The calling into play of these skills and expectations does not explain how one can be in a state of the same kind as that of genuinely perceiving when one is not; it is not meant to do that. What it is meant to explain is the occurrence of a distinct kind of state of consciousness – hallucinatory consciousness – which is such that we can be unable to tell it apart from the genuinely perceptual state. In hallucination, precisely, it seems to us as if we are in contact with situations and things, when we are not. Hallucination *seems* to have *perceptual content*, even though it does not.

Put this way, we can see how we manage to circumvent Martin's worry that we would explain the possibility of hallucination at the price of making the object's involvement in episodes of genuine seeing nonessential. We would have to pay that price if we had managed to demonstrate that hallucinations are, or could be, at least as far as their inward subjectivity is concerned, exactly like states of seeing. But that isn't what we have explained or even tried to explain. The fact that it we may be unable to recognize that we are hallucinating and not actually perceiving does not show that experiencing and

hallucination are states of the same basic kind; it does not even show that they are states of the same qualitative character. In general, the fact that I can't tell a Rembrandt from a Vermeer doesn't show that there is no difference between them. Martin is right that the negative epistemological fact -- that one can't know that one is not seeing -- doesn't itself tell us anything about the intrinsic nature of hallucination; in particular, it doesn't reveal it to be the same as that of perceiving. We have offered an explanation of how it might come to be that we are unable to *tell* one kind of experience (the hallucination) apart from another (an episode of seeing). We have not offered an account of how one kind of experience (hallucinating) might turn out to be of the very same kind as one of seeing. A fortiori, we have not offered an account of how experiences of the same basic kind as perception can come to occur in the absence of the involvement of situations and things themselves. – In the previous section we granted (how could we not?) that perfect hallucination is at least a logical possibility. It is important to appreciate that this can be so even though no perfect hallucination is ever an episode of the very same kind as an episode of genuinely seeing. For the latter, unlike the former, is an achievement of contact with situation and things.

10. Enactive realism

We are now in a position to comment on the so-called disjunctive conception of experience. A main idea of at least a central variant of disjunctivism is that perceiving is a direct relation with situations and things. For this reason, the perception of the wall's yellowness and the corresponding hallucination could not involve the occurrence of one and the same conscious state. The content of the veridical experience, but not the

nonveridical counterpart, is the wall's yellowness. The burden on disjunctivism is that of making intelligible that the two episodes do not share the same content, even though, as far as the subject is concerned, they present or seem to present the world as being the same way. How can the content of perceptual experience depend, in this way, on something external to the perceiver?

I suspect that what makes it difficult to see how the disjunctivist could be right is the fact that we, and perhaps the disjunctivist him or herself, tend to think of experiences as inner representations.¹⁴ If perceptual content is representational content -- that is, if it presents to consciousness ways the world might be -- then it is hard to see how the perception/hallucination pair can fail to share a common content. A perception and its corresponding perfect hallucination disagree, after all, not in how things are presented as being, but in whether things are that way.

We have been considering reasons to give up the idea that experiences are representational in this way. On the view of experience developed here, perceivings are not *about* the world; they are episodes of contact with the world. This standpoint makes it easier to find the disjunctivist alternative intelligible. If perceivings are understood on the model of touchings, then we can think of the content of the experience as given precisely by what we touch. A hallucinatory experience may be indistinguishable from the

¹⁴ I find that there is at least a tension in the work of McDowell on this point. In so far as McDowell holds that perceptual experiences and corresponding thoughts can share one and the same content, then it would seem that he is committed to the idea that perceptual content is representational. But this is a delicate question.

veridical experience, but not because they share the same content, but because it is *as if they do*. A critic may be entitled to demand of the disjunctivist that he or she explain how this *apparent content* is possible. I have tried to do this in the previous section. But what the critic is not entitled to do is insist that the *only way* to explain the indistinguishability of the two experiences is by insisting that they share the same content (or, to put the same point in a different way, that they are of the very same kind). The account presented here seeks to explain the indistinguishability of experiences of seeing and hallucinating without any assumption of shared content (that is, without any assumption of shared contact with the world).¹⁵

The disjunctivist thinks (indeed, *must* think) of perceptual experience according to what John Campbell (2002) calls “the relational conception of experience.” On this conception, perceptual experiences are relations between a perceiver and the environment. Perceptual experience, on this understanding, is not intentional. The sense-datum theorist also thinks of perception as a non-intentional relation to the world, for the

¹⁵ Some resistance to disjunctivism may stem from the worry that the disjunctivist is dogmatic in *disallowing* the demand for an explanation of the hallucinatory appearance of perceptual content. At this late day in history, we tire of being *scolded* for being in the *grip of the Cartesian picture*. In fact, the disjunctivist needn’t adopt the scolding attitude. What the skeptic (or critic of disjunctivism) is not entitled to is the *assumption* that there is no such thing as *the mere appearance of content* or that any *appearance of content* implies that there is content. These assumptions are tantamount to the assumption of the falsehood of disjunctivism.

sense-datum theorist thinks of perception as a mode of direct awareness of sense data. From the fact that one experiences something as yellow, it follows that yellowness exists in one's consciousness. On the relational conception of experience, perceivings are likewise relations to the world in that the content of the perceptual experience of the wall's yellowness is just the wall's yellowness.

Campbell criticizes some formulations of disjunctivism (e.g. McDowell's) for taking intentionality for granted. The world reveals itself to us in perception, and, according to Campbell, it is perception – conscious awareness -- that first makes thought and talk about the world possible. For this reason, thought – even world-involving thought – cannot provide the content of a veridical perceptual experience. Perceptual awareness of things in the world is a more primitive kind of world-directedness than we find in representational thinking.

Understood one way, this seems right. Perceiving is a relation between the perceiver and the world. Perception is non-representational in the sense that perceivings, as I have urged, are not *about* the world, they are episodes of contact with the world. And this contact can surely enable us to think about the world with which we thus make contact. Understood in a different, more foundationalist or empiricist way, however, Campbell's insistence of the primitiveness of conscious awareness strikes me as unjustified. The fact that perceiving is a direct relation to the world – something we cannot do in the world's absence -- does not in itself give us a reason to think that perception is *prior* to thought. For a relational conception of experience leaves open the possibility that only thinkers are able to establish or maintain the right sort of perceptual contact with the world.

In fact, I think, something like this is precisely the case. According to the *enactive* direct realism that I am developing here, there is no perceptual experience of an object that is not dependent on the exercise, by the perceiver, of a special kind of knowledge. Perceptual awareness of objects, for enactive direct realism, is an achievement of the sensorimotor understanding.

We can bring out the immediate relevance of this proposal by considering Campbell's claim that perception enables us to confront an individual substance itself (2002: 116). His example: You hear your neighbors through the wall. What you hear enables you to hypothesize the locations of, e.g. the dish washer and the dining table. You form a conception of that which produces such and such sounds at such and such times, that is, you form a conception of the dishwasher as that which plays a certain functional role as the source of sounds, etc. When you finally get to go next door and actually see the dish washer, you encounter the individual substance itself and so are able to form a conception that is not the conception of an hypothesized object, of something that played a certain functional role. You move from a conception of the object in terms of its dispositions, to a *categorical* conception of the object as the base of those dispositions, a conception of it as the thing that stands before you.

The contrast between a dispositional and a categorical conception of the object is clear; but I question whether this contrast gets marked in perceptual experience as distinct from thought. There can be no doubt that perception enables us to encounter objects as the categorical ground of, say, the way they look or sound. We can be confident that our experience of an apple is not a kind of unconscious inference to that which gives rise to the different appearances of the apple. Experience is not ratiocinative in this way.

Nevertheless, as I have been at pains to argue, these facts don't take anything away from the fact that the apple, as an individual substance, is beyond what we can take in at a glance, from a perspective. In so far as the apple as an individual substance is experienced as present in our experience, it can only be as *present but absent*, that is, as absent, but available. And the nature of the apple's availability is determined by practical knowledge of the sensorimotor relations in which we stand to the apple.

In invoking the role of practical knowledge in this way I am not suggesting that *direct* awareness of the apple as an individual substance is unattainable. To the contrary, the point is that the apple's direct presence to perception can never be more than an awareness of it as appearing this way or that and as available in this or that way. Objects, as categorical grounds of their appearance properties, are given to us *by way* of those appearance properties. There's no avoiding this fact; it is the perceptual predicament.

The heart of my claim is that sensorimotor understanding belongs to the complicated adjustments necessary for attaining and maintaining perceptual contact with the world. The claim is not that the relevant understanding gets exercised in the construction of a representation (a thought, or anything else). Perceiving, on this view, is not representational and so it is not thoughtful in the sense in which thoughts are representations. Sensorimotor understanding, like computational or neural processing in the brain, enters into perceiving not by giving rise to representations of what is seen in us, but by enabling us to occupy a vantage point from which it is possible to see.¹⁶

¹⁶ Compare this with what Campbell (2002: 119) writes: "It is not that the brain is constructing a conscious inner representation whose intrinsic character is independent of

Sensorimotor understanding is the background skill thanks to which we succeed in making direct contact with the world, just as it is the baseball player's skill that enables him to make direct contact with the ball.

This is my picture of the way sensorimotor skills enter into perceptual experience. They are, as it were, skillful ways of achieving contact with situations and things. They are not, as it were, elements of the medium of representation. They are not intermediaries. They characterize our manner of openness to the world (to put the point in the way McDowell might). This conception undergirds my proposed way of resisting the argument from hallucination in the last section. What explains the possibility of perfect hallucination is the fact that it is possible for our sensorimotor skills to get triggered by the wrong thing, or by nothing at all. That is, this is what can explain the fact that might sometimes be unable to discriminate our current situation from one of genuinely perceiving. The view explains the possibility of perfect hallucination with reference to the mere exercise of the relevant skills and it does so, very importantly, without supposing that the exercise of such skills alone suffice to determine the *nature* of our experience, e.g. whether it is perceptual or merely hallucinatory. There is all the difference in the world between the nature of a hallucination and an episode of seeing.

the environment. It is, rather, that there is a kind of complex adjustment that the brain has to undergo, in each context, in order that you can be visually related to the things around you; so that you can see them, in other words." I am making exactly the same kind of point, but extending it to the role of understanding in perceptual experience.

What is at stake in my disagreement with Campbell, really, is a controversy about the *character* of perceptual phenomenology. Just what are perceivers committed to on a naïve realist conception of their perceptual relation to the world? Campbell emphasizes that we think of ourselves, in perception, as confronted by individual substances themselves, as categorical grounds of experience. It is the fact that perception gives or at least seems to give us the world *this way*, that first enables us to form a conception of objects as independent of us. It is certainly no part of our naïve phenomenology that we take ourselves to represent the perceived scene in consciousness. We take the scene – *what we see* – to be out there in the world and to be, in its basic features, independent of us. This is right, but it is only a partial characterization of our naïve phenomenology. Perceptual phenomenology is more complicated than that. This is the point of my extended discussion, at the outset of this paper, of the two-dimensionality of perceptual content. It is no less rock bottom in our phenomenology that we take ourselves, when we make perceptual contact with the world, to do so from a standpoint or viewpoint. When we encounter the world, we do so by encountering how it perceptually appears *from here*. We experience how things are, and we experience how they merely seem to be. It is true that, for example, when we see colors, we take them to be *simple and manifest* in the sense that we naively take them to be *intrinsic properties of surfaces of objects*. But *that* doesn't imply that we take ourselves, when encountering them, to enjoy a wholly unmediated contact with them. Likewise, when we experience the shape of something (a tomato say, which has a rather distinctive ovoid-with-a-furrow shape), we take the shape to be an *intrinsic property of the object*, but we don't take ourselves, in encountering the shape, to be encountering *all of it*, in a way that is not mediated by perspective. When we

see the tomato, we experience its shape, but we experience the shape as *partially hidden* as a result of our vantage point.¹⁷

For this sort of reason, the task for the theory of perception, as I understand it, is to explain the sense in which we are able to encounter the world of mind-independent things out there when we only have ready access to limited bits of things. To explain this, I believe, we need to draw on our understanding. To gain the world as it is apart from us in terms of how the world is given to us, we need to *understand* what we see.

At this point, the enactive approach may seem to come perilously close to conceptualism. According to conceptualism, what enables experience to “hook onto” the world is the perceiver’s conceptual knowledge. Our experience presents us with deer and yellow walls only because we have concepts of those things and so can appreciate what we see as glimpses of them. Conceptualism, however, faces the charge that it overintellectualizes the mind, and moreover, that it makes it mysterious how mere animals can have perceptual experience. The enactive approach agrees with the conceptualist that what enables the barrage of sensory stimulation to rise to the level of perception is our possession of a certain kind of knowledge, but the knowledge in

¹⁷ Chalmers (2005) argues that *primitivism* comes closest to giving a phenomenologically adequate account of visual experience of color. This can’t be right, I think, not if primitivism requires that we take our perceptual encounters with colors to be, as Chalmers says, *edenic*. Even in the Garden of Eden, before tasting of the Tree of Illusion and the Tree of Science, our experiences of color, as our experiences of shape, were encounters with natures that *transcend* the episodes of encounter.

question is not conceptual, it is sensorimotor. The charge of overintellectualizing is thus answered, in two distinct ways. First, sensorimotor knowledge is knowledge of the way sensory stimulation varies as we move; it is knowledge that we share with non-linguistic creatures. Second, sensorimotor knowledge gets brought to bear in experience not in the form of judgment or belief or representation construction. The knowledge is practical, and we *use* it to gain and maintain contact with the world.

The enactive approach, moreover, explains something the conceptualist can't, namely, the distinctively *sensory* way in which, in experience, the world is present.

I have elsewhere proposed that we think of sensorimotor knowledge as a kind of basic conceptual (or *proto-conceptual*) knowledge (Noë 2004). The significance of that proposal is not to assimilate perceiving to the exercise of the intellect, to representational thought. But to assimilate thought and conceptuality to non-representational perceiving. On one way of thinking of object-dependent thoughts, these are representational states that one can only attain by standing in a certain kind of relation to the world. But on another way of thinking about them, they are (like perceivings) episodes of *grappling* with the world itself. Thought, on this approach, is not prior to perception; nor is perception prior to thought. Some thinking is perceptual and some perception is thoughtful.

11. Concluding thoughts: a puzzle about veridicality

In this paper I have proposed that we think of perceptual experiences as temporally extended patterns of engagement with the world, not as things that happen in us. We enact perceptual experience; it doesn't happen to us. Perceptual experiences, then, should

not be thought of as representations, as internal states that are *about* a scene. Rather, they are episodes of contact *with* a scene. This non-representationalist conception of experience squares with perceptual phenomenology, I would say. After all, perceptual experiences don't *feel* like representations. It doesn't seem to us, when we see, as if *what we experience* is represented in our head. Rather, it seems to us as if what we see is *out there* in the world. And it seems to us as if we have a special kind of *access* to what is out there. Our sense of the presence of objects and properties around us, in perceptual experience, is understood in terms of our being skillfully poised to *reach out and grasp* them. Instead of thinking of perceiving on the model of seeing, which is in turn understood on a kind of quasi-photographic or optical-projective model, we should think of perceiving on the model of touching.

Perceptual experiences are not representations, but this fact does not mean that it is inappropriate to speak of their veridicality or nonveridicality. It does mean we need a new model of what this veridicality or nonveridicality can amount to. On a representationalist conception of perceptual content, the question is, in effect, one of matching or satisfaction. How must things be for things to be the way this experience represents them as being? What can we say about the nature of veridicality if we reject the idea that perceptual experiences have representational content?¹⁸

¹⁸ One strategy would be to say that the question of veridicality only arises when *thought* comes into play. Perception cannot be *incorrect* or *nonveridical*; it is just a coming into contact with the world. Perceptions are not like sentences that can either match or fail to match the world. Properly speaking, the question of veridicality, of correctness, does not

To take steps towards such an account, consider the visual experience of stars in the night sky. When you look up in the night sky, you don't actually visually experience *the stars*; what you see, rather, are points of light in the night sky, points of light you reasonably take to be stars (or to be marks or signs or traces of stars). The stars themselves do not enter into your experience. The direct theory of perception fails for the seeing of stars in the night sky. We can see why this is so by considering that, for one's experience of the night sky to be veridical (as we would like to say), there's no requirement that stars be points of light, or that they really look like points of light, or that they be located where we seem to see them. That's why the fact that stars may have gone extinct millions of years ago does not put the lie to our current experience of the night sky.

The case of seeing stars is one where perceptual constancy breaks down. This breakdown introduces nonveridicality. Perceptual constancy also breaks down, for example, when you look down from the height of a very tall skyscraper at cars and people below. The people look to be the size of ants! That is to say, you can't really

arise. But this strategy faces a serious problem. It makes the idea of *correct perception*, of *things being the way they seem*, in some sense only available to creatures capable of representational thought. On this approach, the only way to be mistaken, perceptually, is to come to a false conclusion on the basis of a perceptual encounter. Granted, this sort of linkage between perceiving and thinking exists, but must it exist for there to be *misperception*? I am undecided.

experience *the people* from the top of the skyscraper. After all, there's nothing in the least ant-like about people (in respect of size). That they look the size of ants is intelligible, of course. This means, roughly, that what you see takes up about the same amount of visual field as an ant would when looked at from a normal upright position. What can't be denied is that this is an *incorrect* experience of the people. For what *would* make such an experience veridical? The actual presence of ant-sized people!

These cases of breakdown in perceptual constancy can be contrasted with cases where there is no such breakdown. Consider our example of the silver dollar seen from an angle. You do experience the coin's circularity (or so I have argued) even though the coin looks elliptical from your viewpoint (or so I have argued). Despite the discrepancy between the coin's shape and the way it looks, there's no element of nonveridicality in your experience of the coin. Elliptical is just the way circular coins look when seen from an angle. Likewise, when you experience a wall as uniformly colored, even though there are visible variations in brightness across the surface of the wall, as a result of uneven illumination, you do not *misperceive* the wall's color.

But now the puzzle begins to take shape. What licenses us to say that there is no misperception in the case of the coin and the wall, but there is in the case of the stars and the people down below? In both cases, you might insist, the very same sorts of causal processes join us to distal events and objects. If *elliptical* is just the way circular coins look from an angle, and if *variegated in brightness* is just the way a uniformly colored wall looks when light falls on it in this way, then why can we not say that *ant-like* is just the way people look from this height, or *point-of-light-like* is just the way stars look from

this distance? How can we defend against the charge that we are simply prejudiced and so arbitrary in our attitude to large distances (or time scales)?

Traditional thought about perception has been governed by an optical projective conception of seeing that is closely allied to the idea that experiences are representations. If you take this conception of experience for granted as a starting point, then indeed it is impossible to find a principled way of distinguishing between the sort of perceptual failure we experience when looking down on people from great heights and the normal, veridical variations in perspective that we experience when looking at objects from different angles. I would suggest that this is a reason – a further reason – to give up the optical-projective conception of visual experience. In thinking about how we might do this we can come back to our problem about veridicality with which we began.

Consider: I have insisted that perception can be thought of as two-step in so far as we perceive things by perceiving their appearances from a location. It is sometimes useful to think of appearances in terms provided by projective geometry. For example, it helps to explain why a round thing looks precisely *elliptical* when seen from a certain angle. But from this it does not follow that looks are projections or that we should think of experiences as like internal screens on which the outer world is projected. In this paper I have proposed that we think of seeings as like touchings. On this analogy, seeing how things look is a way of coming into contact with them, it is a way of grasping them. This is why the things themselves enter into visual experience. We see – that is to say, we come into contact with – *them*.

Now, we have considered different ways of failing to establish and maintain contact with things in perception. I have argued, for example, that if you don't

understand what you see – in the sensorimotor sense I have laid out – then sensory stimulation caused by an object won't rise to the level of contact. In the context of the present discussion, we come upon another form of *failure* to make or preserve contact. When we look up at the night sky, I argue, we just don't succeed in making contact with the stars although we do succeed in making contact with the lights in the night sky. The stars are just too far away! Ditto for the people down below on the ground. In contrast, for the skillful perceiver, a glimpse of a coin from an angle suffices to enable him or her to grab hold of the coin visually, as it were. Of course if you move the coin far enough away, it ceases to look circular, or elliptical, and becomes little more than a spot. This reveals not that coins seen from such and such a distance look like spots, but rather, that at such and such a distance, one's ability to see the coin – to maintain contact with it – breaks down.

It is relevant that most cases of nonveridical perception are cases of *merely partial* nonveridicality. You misperceive the spoon as bent in water, for it isn't bent. But in thus misperceiving the spoon, you do succeed in seeing the spoon. In exactly the same way, you may fail to visually experience the stars, but not because you are hallucinating. You *do* experience the lights in the night sky. You *are* in contact with them. For example, when you move your eyes away, they go out of view. You modulate your relation to them in this and many other ways. For the vast majority of cases of nonveridical perception, the world *is* at hand, and is present, thus, as a partner in the experience, as content for the experience.

Perfect hallucination (to use Martin's phrase again) is altogether different from ordinary misperception. For in perfect hallucination, the mistake is more radical. It is not

that you misidentify what you are in contact with; it is that you take yourself to be encountering the world (to be having an experience, with content), when you are not. The mistake, then, is that you take yourself to be *in contact* with something, when you are not. And there is no need, as argued earlier in this paper, to explain that possibility by supposing that you are, in such a case, *really* in contact with something else. There need be nothing there, and so no real experience, no real content.

On an enactive, non-representational approach to experience, then, an experience is nonveridical when it has a different content than it seems to have, i.e. when what you contact is not what you might have thought. Perfect hallucination is only misleadingly described as nonveridical perceptual experience. For in such a case there is no genuine experience of the world at all. There only seems to be.

Perception is a mode of encounter with how things are by encountering how they present themselves to a vantage point. Perceptual experiences are ways of coming into contact with the world, not ways of building up or constructing representations of ways things are or might be. Perception is, in this sense, a non-intentional relation to the world. But it is not a brute, external relation. It is one that our brains, and our minds, enable us, with considerable effort, to achieve.

Works Cited

- Austin, J.L. (1962) *Sense and Sensibilia*. Oxford: Clarendon Press.
- Ayer, A. J. (1955) *The Foundations of Empirical Knowledge*. London: Macmillan & Co.
- Broackes, J. (1992) The autonomy of colour. In Charles, D. and Lennon, K. (eds.)
Reduction, Explanation and Realism. Oxford: Oxford University Press.
- Campbell, J. (2002) *Reference and Consciousness*. Oxford: Oxford University Press.
- Chalmers, D. J. (2005) Perception and the fall from Eden. To appear in Szabo Gendler
and Hawthorne (eds.) *Perceptual Experience*. Oxford University Press.
- Dennett, D. C. [2001] Surprise, surprise. [Commentary on O'Regan and Noë].
Behavioral and Brain Sciences 24, 5.
- Dennett, D. C. [2002] How could I be wrong? How wrong could I be? *Journal of
Consciousness Studies* 9, 5-6: 13-16.
- Hurley, S.L. & Noë, A. (2003) Neural plasticity and consciousness. *Biology and
Philosophy*, 18: 131-168,
- Kelly, S. D. (2004) On seeing things in Merleau-Ponty. Carmon, T. (ed.) *Cambridge
Companion to Merleau-Ponty*. Cambridge: Cambridge University Press.
- Merleau-Ponty, M. (1945/1962) *The Phenomenology of Perception*. Trans. Colin Smith.
London: Routledge Press. [French original: 1945].
- Noë, A. (2004) *Action in Perception*. Cambridge, MA: The MIT Press.
- Noë, A. (2005) Experience without the head. To appear in Szabo Gendler and Hawthorne
(eds.) *Perceptual Experience*. Oxford University Press.

Noë, A. (2002) Is the visual world a grand illusion? *Journal of Consciousness Studies*.

Volume 9, No 5/6: 1-12.

O'Regan, J.K. and Noë, A. (2001) A sensorimotor approach to vision and visual

consciousness. *Behavioral and Brain Sciences* 24/5: 883-975.

von Senden, M. (1932/1960) *Space and Sight*. Trans. Peter Heath. London: Methuen.

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

Putnam, H. (1999) *The Threefold Cord*. New York: Columbia University Press.

Smith, A.D. (2000) *The Problem of Perception*. Cambridge, MA: Harvard University Press.

Valvo, A. (1971) *Sight Restoration after Long-Term Blindness: the Problems and Behavior Patterns of Visual Rehabilitation*. New York: American Federation for the Blind.