

Dear Prof. Byrne,

I have been interested in philosophy – especially the question of the role of consciousness in the physical universe – since I was young and had a kind of epiphany about what seeing involves (more of that below), but I am not trained as a philosopher. So given your vast reading in the philosophy of color, and your recognized work in this area, it might seem that I could have nothing to offer you that could improve your understanding of your chosen subject.

Normally that's what I would have thought. I am used to reading papers in philosophy – especially about direct perception – which I find incomprehensible. This seems to be a failure on my part to imagine what the author could be thinking. "How on earth can they believe such things?" I am apt to ask myself, because I cannot reconstruct the assumptions that allow them to make the assertions they do, or perhaps because my own assumptions prevent me from comprehending theirs. Clearly this does not put me in a position to teach anyone anything.

However, when I read your chapter in *The Oxford Handbook of the Philosophy of Mind*, "Sensory Qualities, Sensible Qualities, Sensational Qualities", I felt that I was in the opposite position: it was you who were failing to comprehend a position which I hold and which I find obvious. (My position is one which allows me to read passages on perception by Hume and Locke, as well as the quotes you critique in your chapter, and feel that they are just expressing what seems to me to be common sense, or what should be common sense to the scientifically sophisticated: perception is indirect. We are immediately aware of mental representations, not the outside world. Sensory qualities are properties of these representations.) I found this apparent misunderstanding on your part to be quite amazing, given your obvious expertise in the field. This reversal of positions seems to give me something of value to teach you, if I am able. Even if, once you understand the position you are critiquing as I would have you do, you can muster strong arguments against it (and I would be very interested to know those arguments), you cannot do that until you see the target clearly. Several times in the paper you say "it is unclear", "it is not immediately clear" as if this were a criticism instead of an admission of incomprehension. I'm sure you will agree that incomprehension is not a valid form of argument. So let me try to clear a few things up, if I can, by explaining what I would have meant if I had written some of the passages you analyze.

A major source of confusion in your paper, in my view, is your assumption that the claims of indirect realism are meant to be "pre-theoretical". You say of one of your unclarities that it "casts some doubt on the idea that 'sensory quality' is a label for something we all pre-theoretically recognize." But the pretheoretical position on vision is naïve realism, just as geocentrism is the pretheoretical view of the sun's motion through the sky. If a Copernican were to claim that when we say "rise" in "the sun rises" we don't mean it in the same way as in "smoke rises", and someone who dismissed Copernicus's theory out of hand, to the point that he hadn't even attempted to imagine what it would be like to live in a Copernican world, or what a Copernican might mean when he said certain things – if such a one criticized the Copernican's statement on linguistic grounds, or because the distinction wasn't clear to him, he

Sunday, October 25, 2015

would be missing the point, because when the Copernican said "we don't mean" he would have been referring, not to "the vulgar" (as Hume would say) but to those who have embraced the radically revisionist theory that the earth moves. The distinction he intends can only be understood within that theoretical framework. To refuse to recognize the theoretical assumptions a statement assumes is just to invite confusion. And that, it seems to me, is what you do throughout your chapter.

Indirect realism was just as radical a revision of the common sense perspective as was Copernicanism, and took hold in the same early modern period, soon after Kepler gave us the modern understanding of the optics of vision: light comes into the eye, and is focused by a lens onto the retina into an image of the outside world. Contrary to an ancient and widely held view, nothing goes out from the eye, allowing the mind to touch things where they stand. (Studies by G. A. Winer et al. have shown that even today, if asked, a large proportion of even educated adults – including students studying perceptual psychology – will volunteer some version of the latter theory. Naïve realism fits easily with, and may even suggest, this age-old "extramission" theory, but is less easily reconciled with the true physics.) I think many early modern natural philosophers, on the basis of this new knowledge (at least in part), must have gone through a kind of sudden realization about the nature of vision very like the one I did when I was a boy.

I was somewhere between 10 and 12 years of age and very interested in science when I had a kind of epiphany. At least that's the way I remember it. I had been puzzling over perception, and was looking at my mother's rose garden, when I imagined a kind of darkness there instead, because it was a world without color. It was made of atoms reflecting radiation of different frequencies. These rays differed in number, not quality. They reflected into my eyes, which sent nerve impulses to my brain, somehow causing the colored world of appearance to be created. It was not out there; it was inside me, like a dream or a movie. But mainly I remember this impression: as if standing on the edge of an abyss, looking with my mind into incomprehensible darkness, into a world without sense. That day the world, as I liked to say, turned outside in and inside out. And it has never turned back. Not that I constantly feel like I'm in a dream or that the world is alien and unknowable, but I understand that this sensual world is subjective, and that the world outside, which causes my brain to create this appearance, will always be, in a sense, beyond my reach, beyond "the veil of perception."

Unfortunately, my notebook from that time is lost, and I don't remember my immediately preceding thoughts, or any words that went along with the realization, but I am certain that the so-called "argument from illusion" played no role for me. I was not concerned with the possibility of misinformation (optical illusions or hallucinations) but with the form the information took. It was the profound difference which I now believed existed between appearance and reality which struck me with wonder. I do remember trying to tell people afterward that the quality of redness is in the mind, that it is a quality of consciousness. No one seemed to understand, or be interested, though it seemed to me that the greatest mystery in the universe was how redness came to be from something like electrons moving around in the brain.

Sunday, October 25, 2015

Later what struck me as particularly paradoxical about the naive view, given the directionality of the visual process (it is all one way: inward), was how anyone could believe that the final product of such a process could be the world outside us, itself, appearing to us where it is, at a distance. I wanted to ask, "How did it get back out there?" It seemed obvious to me – and still does – that the world's appearing, at the end of this inward-bound process, must be taking place inside the skull, even though that appearance is ABOUT what is outside us. But just exactly how that argument is supposed to go I'm still not sure.

Sometimes I wonder whether the division in philosophy between direct and indirect realists is between those who have had such an epiphany and those who haven't. Or perhaps there is a second epiphany, which switches everything back. I have sought it in vain. Direct realism makes no sense to me.

If "epiphany" is apt, then our differences are rooted at least as much in imagination as arguments. I cannot imagine a direct realist world consistent with what I know of science, and perhaps you have not imagined living in an indirect realist one. That, at least, is what I believe your arguments indicate. How else could you say something to me so absurd as, "there do not seem to be any mental candidates for the smell of the lilacs: if the smell is anywhere, it is in the lilac bed, not the mind"? Your discussion of twinges also seems to me to reveal a complete failure to imagine how indirect realists construe the world. You say,

I am aware of a twinge, its location (my left elbow), and its achiness. That sort of awareness seems (quasi-)perceptual—rather like the awareness I have of a flash of light, its location (the far side of the room), and its brightness. However, on this conception of (bodily) sensations, they are not experiences—rather, they are objects of experiences. The twinge is—or at any rate appears to be—some sort of disturbance in my elbow, not in my mind or head, unlike the presumed 'experience' of the twinge. Clark, however, has a very different conception of bodily sensations, on which the experience of the twinge, not the twinge, counts as a bodily sensation....

What a telling and unfortunate ambiguity: "The twinge is—or at any rate appears to be—some sort of disturbance in my elbow, not in my mind or head, unlike the presumed 'experience' of the twinge." Did you mean "unlike the presumed 'experience' of the twinge which, if it existed, would be in my mind or head, not in my elbow," or "which would appear to be in my mind or head, not in my elbow", or both? You can say "is—or at any rate appears" because you don't posit a radical difference between appearance and reality. For Clark and me, appearing to be in the elbow is a very different matter from being there, and being in the mind or head does not imply appearing to be there. Indeed, that's the whole point of appearance, i.e. the immediate object of perception: it appears to be out there, but it is all in the mind. For us, there is only one twinge. It is an experience in the mind felt to be located in the experience of the elbow, which is also in the mind, caused by a physical disturbance (not a twinge) outside the mind in the physical elbow.

Sunday, October 25, 2015

Assume indirect realism is a theory, or rather a class of theories, within which sensory, sensible and/or sensational qualities are supposed to play a role. Here are several analogies which express aspects of that theory:

1) Waking perception is like experiencing a very vivid dream. But while dreams are created by the free play of the mind, these waking dreams are under the control of sensory input.

2) James Cameron, in directing "Avatar", used an invention called a "virtual camera" or Simulcam. As his actors performed, surrounded by video cameras, he held a device with a screen which showed him, in real time, the output of computations based on the actors' motions and the position of the screen. The actors appeared to him as his blue creatures, the Na'vi, situated in an alien landscape. As he moved the device, the point of view of the computed image changed, just as if he were moving a video camera through the space inhabited by the Na'vi. He was experiencing a virtual world based on computation and optical input. So are we.

The computational machinery responsible for our 'display' reveals itself when the machinery itself is overloaded or fatigued (afterimages, motion after-effects) or its algorithms have shortcomings (Müller-Lyer illusion), or when it seriously malfunctions (hallucinations), or perhaps, under exigent circumstances, when it seeks to communicate something important to the 'user' via visions and/or voices of supernatural beings.

Not only do we directly experience a virtual world computed in near-real time based on sensory input from the real one, our bodies participate in both. It's as if Cameron, wearing head-mounted cameras, could see his own body as a Na'vi, from a Na'vi's perspective, immersed in the Na'vi world. That blue figure would become the visual version of his body image. (A body image is a mental representation of one's body, used for both perception and control.) The fact that sensory-based computed versions of our bodies appear and act in our visual, tactile and sensorimotor spaces allows these spaces to be coordinated and merged with each other (as the sensory versions of our body images are), and to take on the functional attributes of physical space. This is how you can look at what I point to. This is how we share a world.

When you feel a pain in your knee, it is caused by injury (not pain) in your physical knee. You feel the pain located in your body image's knee – it appears to be there in what appears to be your knee. When you point to the pain, your body image's hand points to your body image's knee, your physical hand follows suit, pointing to your physical knee, and your friend is informed where your painful injury lies.

3) 3D TV. Imagine you are sitting in a dark room with a bright rectangle on the far wall. In that rectangle you can view a ball game as seen from behind first base. You are definitely seeing the game. But are you seeing it through a window or on a 3D television screen? There are some knobs by your right hand. As you turn the knobs you find you can change the color saturation and brightness; you can distort the image, making it seem to flow in unrealistic ways, just as we can do to our own visual experiences by taking drugs; you can make the ball field seem to move

as it would if you were seasick or dizzy. You decide you are seeing the ball game on television (indirectly).

Were you to decide instead that the TV screen is a window, and that the video images are the players themselves, but you were wrong, this would be a kind of illusion in itself, even if all the information conveyed by the images were accurate.

Of course, if you had been told that the game was actually being played in a distant city, and that information about it was being transmitted to your room first as an analog signal by optical cable and then encoded digitally, you would not have had to touch those knobs to know that you were viewing the game indirectly, even if you had no idea how the display was created.

4) When astronomers view digital images from their telescopes they often use pseudo-colors. These may represent intensity or wavelength in X-ray or infrared, types of radiation that are invisible to us and so have no color. Earth imaging uses pseudo-colors to represent all kinds of features: temperature, type of crop, elevation. Petroleum engineers might view the strata beneath their feet, where no light shines, in pseudo-colors. Soldiers view infrared as glowing green. Political operatives make maps of blue and red. For indirect realism, all colors are pseudo-colors. They are assigned by our biology to mentally represent certain optical reflectance properties in the outside world to which our eyes happen to be sensitive.

5) In the movie "Terminator" there are scenes in which you see as if 'through the eyes' of the cyborg. Image outlines are annotated alphanumerically. We can imagine that our vision could have been like that. Each part of an image would be tagged with a string of numbers representing its spectral properties. The numbers would be assigned by unconscious computations and available for use by conscious ones. Instead what we do have is parts of the visual field color-coded (with pseudo-colors) to represent spectral properties. But the principle is the same. Neither kind of tag is in the outside world; they only represent information about that world. Such 'tags' are not "raw feels" – they are not the raw data upon which the representation is computed; they belong to the display.

I would go further: just as the Terminator annotated images with all kinds of technical data, there are many other kinds of 'tags' which are computed and which are parts of the experiences we have of objects, although they are not typed out next to them or coded by colors: face, tree, tomato, car, table, beautiful, sexy, sad, the letter 'b', the word 'Stop'. Like colors, these are also presented to us integrated into the perception as properties of the objects we see. In James Cameron's case, the computer couldn't output the appearance of a Na'vi unless it had already determined to interpret a particular set of inputs as a Na'vi. We get the already-interpreted outputs, pre-identified and appropriately structured (into, say, the drawing of a rabbit or a duck). That's what we say we see – not a "bulgy shape" but a tomato or a face. We always see "as". The visual world comes already endowed with meaning; the same for the other senses. For that reason, I don't believe "sense-data" is a good name for the immediate objects of perception.

6) C. D. Broad wrote (in *Scientific Thought* 1959/1923, p. 247), "The best analogy that we can offer to the relation between our sensing of a sensum and our perceiving a physical object, is to be found in the case of reading a book in a familiar language. What interests us as a rule is the meaning of the printed words, and not the peculiarities of the print. We do not explicitly notice the latter, unless there be something markedly wrong with it, such as a letter upside down. Nevertheless, if there were no print we should cognize no meaning, and if the print were different in certain specific ways we should cognize a different meaning. We *can* attend to the print itself if we choose, as in proof-reading. In exactly the same way, we are not as a rule interested in *sensa*, as such, but only in what we think they can tell us about physical objects, which alone can help or hurt us." On this analogy, the theory of "transparency" (what Broad (p. 246) called "a baseless prejudice ... that we are not aware of *sensa* and their properties") is like claiming that we read without seeing the page.

In your chapter "Skepticism about the Internal World" (which I found recently, a year after writing most of this letter, but then setting it aside) you take up Broad's metaphor, but draw the opposite conclusion. You write,

"The point is that the newspaper example is *not* a good model for how you know you see something. That is, if I ask you "Do you see a cat?" you do *not* have to attend to something that is analogous to the newspaper (perhaps a "sensory experience" or a "visual sensation"). When you read about the Red Sox, you don't just find facts about baseball, you also find the newspaper. But when you open your eyes, your seeing is in a way invisible. What you initially find is the *world*, not your *seeing* of the world."

If this is meant to be an argument, it begs the question. How do you know that you don't have to attend to something analogous to the newspaper, namely your visual experience? I would say that you are attending to it (as opposed to attending, say, to your tactile experience), and this allows you to attend to the cat, by attending to its visual representation. What you don't need to do is *know* that you are attending to your visual experience, as opposed to the world that it represents. Normally we think of them as one and the same, and this is good enough for all practical purposes. You say, "When you read about the Red Sox, you don't just find facts about baseball, you also find the newspaper." And when you see a cat, you don't just find facts about the cat, like a list of propositions, you find the way the cat appears to you, and from that you can read off facts about the cat.

Perhaps this will help (though I doubt it): The cat and the appearance to you of the cat are not the same. When you close your eyes, one ceases to exist and the other doesn't.

To answer a few obvious objections:

I don't mean to imply that there is a "little man" inside the head viewing an image. Obviously, that would lead to an infinite regress. This stumped me for a long time, until I came to think of the "stream of consciousness" as *consisting* of our conscious experiences. These experiences are not being viewed from outside themselves. They don't need to be re-presented; they are the

representation. Another way to put this is that appearance is experience. I don't believe that there is a so-called "act of apprehension" separate from the appearance apprehended. To appear is to be apprehended (which is not to say that focal attention is trained on it.) To use the model of seeing here is a mistake. What we see – the object in the external world – is always apart from us, and its appearance appears to be at a distance from us. It is natural to model our understanding of our relation to that appearance on our relation to the distal object itself, that is, to model apprehending on seeing. Seeing is a relation between a subject and a distant object, so apprehending the appearance of that object is taken to be a relation between a subject (never seen but implied by the point of view) and the appearance of the distant object. But appearance (as I conceive it) is inherently mental. It is already an act of mind, without requiring any further act to appear. In apprehending an appearance there is no separation between subject and object in need of bridging by such an act. If this is true, there is no "experience", conceived as separate from its immediate object, to be "transparent" with respect to it, as some claim. Transparency, under this theory, is otiose. Moreover, if appearance is so utterly unlike the external world as this theory contemplates, the idea that perception is "transparent" takes on a comical air. It is as if one watched a French movie dubbed in English and exclaimed, "The French comes through so clearly it's as if I was standing right there among the actors!"

The objection that when we see green there is nothing green in the brain is based on the direct-realist assumption that green as we experience it is a real property in the outside world. But if things we see as green are those which reflect light in such a way that it causes humans to have an experience of green, there is no reason to believe that, just because we are experiencing green, there is anything in the brain with those reflectance properties.

Likewise for spatial properties: there need be nothing square or huge in the brain when we see something square or huge. The dimension experienced in vision as length is not the dimension measured in the outside world by yardsticks, anymore than the auditory experience of pitch is identical to wavelength. Mental space is not physical space. But that does not mean, I hope, that it is not physical in some sense, or that it is not located in physical space in the brain. To decide this, one would have to go into what it means to be physical and to be in a location – and of course, solving the mind/body problem would help.

Even with its many unknowns, the basic claims of this theory seem perfectly intelligible to me. If you have a knock-down argument against it, I would like to hear it. Or if you have other kinds of reasons for rejecting it, I'd like to know what they are.

From the point of view of this theory, as I've mentioned, naïve realists mistake appearance for reality. They take the mental representations of external physical objects to be the objects themselves. They take the properties of these experiences (for our mental representations *are* our experiences) to be physical properties in the outside world. Every child is a naïve realist, and so are most adults, so all our everyday ways of speaking are based on this illusion. It is no accident that the illusion is so compelling. The whole point of having an internal representation of the world is so we can interact with the world on its basis. Evolution produced it for this

reason. But evolution had no reason to inform us that the mental representations we are using to interact with the world are not the world itself. This only became clear (to a few of us) when we found out how vision works. And even so, it is hard to keep in mind. — At least that is how things are according to my version of the theory of indirect perception. I don't claim that any of the philosophers you critique hold the theory in just this form, but I believe the accounts line up enough for me to understand without difficulty the passages you analyze.

In your paper you use this quote from Levine:

Let's take my current visual experience as I gaze upon my red diskette case, lying by my side on the computer table. I am having an experience with a complex qualitative character, one component of which is the color I perceive. Let's dub this aspect of my experience its "reddish" character ... Qualitative character concerns the "what" it's like for me: reddish or greenish, painful or pleasurable, and the like. From within the subjective point of view I am presented with these qualitative features of experience, or "qualia," [a.k.a. sensory qualities] as they are called in the literature. Reddishness, for instance, is a feature of my experience when I look at my red diskette case. It is notoriously difficult to explain this feature by reference to either the physical or formal features of my brain states.

This passage is completely consistent with the theory I have just painted. But you proceed, in my opinion, to completely misunderstand it. You say (and I comment in brackets),

Sensory qualities—if there are any—are properties of experiences. They are not properties of external objects like diskette cases and the like. [So far so good.] Going by the quotation from Levine, sensory qualities are not even ostensible properties of diskette cases—properties of experiences that appear mistakenly to be properties of diskette cases. [Wrong!] Introspection, attention to one's experience, acquaintance, or something similar, seems to reveal them as properties of experiences [No.], and that is indeed what they are [True].

Ostensible properties of diskette cases – properties of experiences which are mistakenly taken to be physical properties of the objects of those experiences by naïve realists – are *exactly* what sensory qualities are. It is not introspection which reveals these ostensibly physical properties to be properties of my experience. It is the realization that they must be, given a science-based understanding of the process of vision (that it is a one-way causal chain going inward, conveying information, not qualities – differences, not what the differences are between – so the qualities themselves must be provided by us). That is, it is the theory of indirect perception which causes me to reinterpret my visual experience. It is not a matter of noticing, by introspection, anything *in addition* to what I noticed before. Rather, if I believe that seeing is like having a vivid dream based on sensory input which conveys information about a world whose true nature is beyond imagination, then I can conclude that all the qualities of this dream must be subjective; all colors are pseudo-colors; all perception is introspection. I have faith, because I am a realist, that there

is an isomorphism between my experiences of the world and what is out there, but that is the extent of the similarity I can count on.

Once this split occurs, between appearance and outside reality, to avoid confusion one needs to make a clear distinction between physical properties and subjective ones. That, I believe, is why Levine introduces reddishness: "Reddishness ... is a feature of my experience when I look at my red diskette case." To properly understand this sentence you must assume the theory of indirect perception. By "red" Levine means here the physical property, outside of his experience, which causes the object in his 'waking dream' which represents the diskette case to be experienced as "reddish". That very same quality (reddish) is what a naïve realist calls "red".

At one point you do acknowledge this distinction, discussing it as an ambiguity: "in one sense, 'red' refers to a property of tomatoes and strawberries; in another sense, 'red' refers to a property of mental items—experiences, sensations, or whatever." However, you object to this "odd talk": "taken as a descriptive semantic claim about English, the ambiguity thesis is not very appealing." But, as I have said, everyday language assumes a naïve realist perspective, just as "the sun rises" assumes a pre-Copernican one. These facts do not constitute evidence against science-based, revisionist theories.

But I would go further. There is nothing odd about talk incorporating this kind of ambiguity. Later in the chapter, you advance this clever argument against G. E. Moore's claim that "'blue' has 'two different senses' when we speak of a 'blue afterimage' and a 'blue tie'":

... there appears to be no ambiguity: 'My afterimage and this tie are both exactly the same shade of blue' would strike us as false if 'blue' were ambiguous between 'the color blue' (not a property of afterimages) and 'blue [prime]' (not a property of ties). On the contrary, this sentence seems unproblematically true.

Your conclusion seems initially plausible, but consider this parallel case:

Two political operatives are discussing American politics. The familiar political map of the country hangs on the wall, with predominantly Republican states colored red, Democratic ones blue, and those in the middle various shades of purple. New data have come in. John asks Joe, "What does Kentucky look like?" Joe answers, "It's exactly the same shade of purple as your tie."

Joe's answer seems to me to be an entirely natural way of talking. In the context of their representational system, Kentucky is purple. By calling it purple, I mean that it has a political complexion that would be well represented by that color, in contrast to say Mississippi or Massachusetts. (States don't come any redder than Mississippi, while Massachusetts is the deepest blue.) John's tie has no political complexion, yet it can share a shade with a state, which (in the political sense implied) has no reflectance properties. Likewise, within the context of our perceptual representational system, a physical surface with certain reflectance properties can

be the same color as an afterimage in the sense that it is normally represented subjectively by the very same sensory quality as the one the afterimage exhibits.

After writing the purple tie example above, I actually heard this on the radio: "Oklahoma is among the reddest of red states, as red as the meat served at the Gentlemen's Steak House." (Brian Nailer on All Things Considered, NPR's evening news program, June 23, 2014)

It may be tempting to claim that the primary meaning of 'red' is either the subjective quality or the physical one. But as this example makes clear, each has its role to play. The objective (represented) properties out there in the world are what we need to know about. The subjective (representing) ones allow us to perceive, think and speak about them. It is a testament to the flexibility of our cognitive systems that we can also talk about the subjective qualities themselves, and that we are capable of making the distinction. But for some this is evidently very difficult.

By refusing to make this distinction you double qualities in a way not contemplated by the theory – attributing to it claims it would never make – thereby sowing confusion. You say, "Levine's sensory quality of reddishness, then, is not redness, the color." This phrase "redness, the color" is fatal. As young naïve realists we learn to name the colors. For us "red" had only one meaning: it was the quality we experienced when we saw red things, a quality we took to be out there in the sensuous physical world. That experienced quality which we called red is exactly the quality Levine is calling "reddishness" to distinguish it from a quality which we, as children, never imagined: the set of physical properties out there in the world which have no necessary relationship with our experience of the color red (that is, with reddishness), but which are its normal causal antecedents. This set of physical properties is what Levine assigns the word "red" to. In this theoretical context, the phrase "redness, the color" has no clear meaning.

To continue from the sentence of yours I just quoted:

Levine's sensory quality of reddishness, then, is not redness, the color. Even if tomatoes are not in fact red, at least they appear red. And one's experiences do not appear red at all. Levine himself is quite explicit that redness and reddishness are distinct (see Levine 2001, p. 179, n. 5).

On the contrary, tomatoes as they appear in one's 'waking dream', which *are* "one's experiences" (of tomatoes) – do "appear red" (in everyday language), or, in Levine's technical vocabulary, they have the quality of reddishness, which causes us (direct and indirect realists alike) to believe the tomatoes to be (physically) red. This is the distinction which Levine makes in the note you refer to, namely that between "features of physical surfaces" and "features of experiences". He does not distinguish "reddish" from "the experience of red" or "the appearance of red" because these are one and the same. He only distinguishes it from red, the physical property.

(At least this is my reading of the relevant passages from Levine. What I've really been talking about are not Levine's true positions, which I haven't read enough of to even claim to comprehend, but what I would mean if I had written the relevant passages. I apologize to Levine if I have misread him.)

This understanding (or misunderstanding) of yours, that reddish experiences would not appear red, turns out to figure centrally in what you call "the decisive defect" of "natural sign theory": that it "cannot accommodate the plain fact that objects look colored."

Here's your account of the natural sign theory:

Imagine that someone is looking at a tomato, and sees that it is red. If we apply the basic outline of the natural sign theory to the perception of the tomato's color, what is going on? First, a property of the tomato causes a certain effect in the perceiver: a reddish sensation. Second, she becomes aware that this sensation is reddish. Third, she forms the belief that the tomato has the power to cause reddish sensations (i.e. she believes that the tomato is red). End of story.

But that's not the end of your story. You continue:

Of course, the tomato will look red to the perceiver. The natural sign theorist can say that this simply amounts to the fact that the perceiver believes that the tomato is red and has a reddish sensation.

You then call into question whether belief plus reddish sensation could result in the tomato's looking red. For instance, according to this account, a failure to believe the tomato is really red would result in a failure of the tomato to appear red, which does not happen.

As far as I know, your analysis of "appearances ... in terms of dispositions to believe" is a straw man wholly of your own making. Reid speaks of color sensations, which for him are "ideas", as "appearances" of colors. So the way something appears comes first for Reid. It does not rely on belief or inference to unknown causes (although judgements made on the basis of appearances may.) That is the case for me also, and for the other philosophers you discuss. Your distinction between perception and appearance is not ours. So the "decisive defect" you attribute to indirect realism belongs only to your misinterpretation of it.

As I mentioned at the beginning of this letter, I often find things that direct realists say incomprehensible, no doubt just as incomprehensible as you seem to find the opinions I share with the philosophers you critique. Your ability and inclination to doubt the very existence of the mental entities which I take to make up the fabric of my world is beyond my ability to comprehend, at least so far. I don't know how to entertain that radical doubt of my theory of perception because I don't understand at all what you want me to put in its place – or, to the extent I do, it doesn't make any sense to me.

Sunday, October 25, 2015

Your recent chapter "Skepticism about the Internal World" would seem to be an excellent chance for me to learn from you how to see things your way, but unfortunately I found it, like your other writings, as maddening and perplexing as ever.

This letter is already far too long, so I will only mention one problem I had with your "Skepticism" chapter, your misuse of the word "evidence": you say "evidence about" when you should say "evidence provided by". For instance, you say (with my underlining added),

"Now consider certain non-human animals, for instance dogs. They have sensory experiences (or so we may suppose) but there is not much reason to think that they *know* that they have sensory experiences. Knowledge of one's own mind requires a sophistication that dogs appear to lack. So if a dog knows that there is a rabbit behind a tree by using its eyes, it is not on the basis of evidence about its sensory experiences. And if evidence about sensory experiences is not needed for a *dog* to have environmental knowledge, it isn't needed for *us* to have environmental knowledge either. You could know that you are sitting in a chair without appealing to evidence about your sensory experiences, and presumably you do."

I don't know I am sitting in a chair on the basis of any evidence about my sensory experience (whatever that could be), I know it on the basis of the evidence provided by my sensory experience, that is, I know I am sitting in a chair on the basis of my sensory experience. This does not require higher order knowledge or belief about my sensory experience... unless someone ups the ante and challenges my belief with skeptical doubt. Then I will need to provide evidence for trusting my sensory experience – evidence about it. This, of course, dogs cannot do. But dogs and naïve realists can know they are sitting in a chair on the basis of their sensory experience simply by taking aspects of that experience to *be* the chair, and other aspects of it to be their bodies, without knowing they have minds at all. It seems to work fine – as practical knowledge, although based on an error. But of course it's always possible that they are only dreaming.

This last example seems to follow a pattern of yours, where you attribute to your opponent's theory a claim of higher order awareness than the one intended, and then dismiss it when the implications of that self-awareness are falsified: Dogs must know *that* they have sensory experience, not merely know their sensory experiences. To see that a tomato is red requires a person to become aware *that* her sensation is reddish. Appearance is based on belief. Evidence *about* one's sensory experiences is supposed to be required for sensory experiences to serve as the basis of perception, instead of just the sensory experiences themselves serving as evidence. For Clark, supposedly "the experience of the twinge, not the twinge, counts as a bodily sensation." (No, as I explained, the twinge is the bodily sensation, which is a kind of experience, and it is felt to be located in the body. The object of a twinge is a physical injury or disturbance.)

I don't think these are intentional misrepresentations. I think they result from your failure to get an imaginative grasp of indirect realism's view of what's going on in perception. You seem to think that "visual experience" is meant to be something in addition to what you are already

aware of when you see, and when you can't find this extra something or its purported properties you doubt their existence and reject indirect realism. If you truly understood, I don't see how you could say,

"Suppose that Hume and Hayek are right, and that nothing in the external world is colored or noisy or tasty. What remains quite unclear is why they think items in the 'internal' world—sensations, for instance—have these qualities. Why not say instead that absolutely nothing has these qualities?"

This makes absolutely no sense to me. If nothing had these qualities, we could not experience them. And we do, every time we experience anything.

I sympathize with your difficulties, because I completely fail to get an imaginative grasp of what *you* believe is going on in perception. You say you suppose dogs have visual experiences, but you imply that they don't rely on them in order to see. What is the role, then, of visual experience in vision? Can we see without it? Do we? I don't understand. How do *you* know you are sitting in a chair?

I really wonder, assuming you have read this far, whether what I've said will have had the power to change your understanding of the indirect realist philosophers you critique in your various articles. I've done my best. If, in your view, it is I who have misunderstood you, then I apologize, and am eager to be corrected.

I hope by now you understand what I mean by 'appearance'. I interpret it as something mental, but regardless of how we conceptualize it, when we see, the world appears to us. Would you agree? What I need to know from you, if I am ever going to understand your position, is this: Is there a place for appearance in your theory of perception? (If there isn't, I don't see how I could believe it.) If so, how does it happen? Where is it (if that makes sense)? If the appearance that we are treated to when we see the world is the world itself, that must mean that the appearance is at a distance from us. Is that right? Finally, what is its relationship to the brain?

At least, those are the questions I would ask you if we were friends talking over a beer. But you are a busy professional philosopher, a professor, author, editor, and chair of a department, and I am a struggling amateur, and a stranger. So I will ask you instead if you could kindly steer me to some references that might be able to cure my perplexity.

I hope I have not offended. Any response at all would be welcome.

Best regards,
Gerald Lame
San Diego, California

Sunday, October 25, 2015