

Commentary on David Rosenthal, "Mental Qualities Without Consciousness"
Pacific APA, April 17, 2014

I want to divide up my comments into three sections, corresponding to what I see as three different issues dealt with in Rosenthal's presentation. They are of course interrelated, but nevertheless I think they merit separate attention.

Methodology and the Status of Intuitions

Rosenthal's target is the allegedly common sense claim that the mental qualities involved in conscious perceptual/sensory experiences are essentially conscious. Consider an ordinary conscious visual experience of a ripe tomato, and focus on the redness that is experienced. Could that very quality - that redness - exist, or be instantiated in an unconscious mental state? Some, including myself, claim that it couldn't. Being a constituent of a conscious experience is essential to such qualities. But what is the status of such a claim, what justifies it?

Many people would say that it's just a fundamental intuition. Now "intuitions" have come in for a lot of abuse of late as a source of knowledge, or a constraint on philosophical theories. Rosenthal quotes Kripke to the effect that intuition seems to be the most sound basis for constraining theorizing, but many others would side with Rosenthal that intuition is itself untrustworthy. In particular, Rosenthal argues, what we think of as simple intuitions are often really the consequences of unarticulated theories. If that's right, then the theories that underlie the intuitions must be brought to light and compared with rival theories. With regard to the intuition in question - that mental qualities are essentially conscious - Rosenthal argues that his rival QS theory actually is superior to the one underlying the intuition, and so therefore the intuition should not be taken to be definitive in the way some are inclined to do.

I can't take on here the entire debate over the epistemic status of intuitions, or their role in justifying philosophical theories. Let me just briefly say how I see the issue, without attempting to argue against all the competing views. First of all, unlike some, I do not take intuitively based judgments to reflect any special cognitive capacity. I do not adhere to the view that we have a faculty of rational intuition (though how we know mathematical truths is beyond me). Rather, I see intuitive judgments as reflections of all-things-considered cognitive states. That is, when I judge that, ultimately, this claim just seems more plausible to me than that one, I'm manifesting the resolution of my cognitive state with respect to the claims in question. I can of course reflect on my reasons and criticize them, a process aided by the objections of others. But, in the end, sometimes one side of an issue just seems the most plausible. While, on the one hand, I think Kripke's right that, in a sense, what more than how things ultimately seem to me could be more worthy of basing a judgment on (I should believe what seems most plausible to you instead?), I also agree with Rosenthal that a lot of unarticulated theory could be behind such judgments and they must always be treated as revisable in the light of future considerations. But of course, these considerations have to convince me in the end, which I take to be Kripke's point.

That's what I want to say about intuitive judgments. However, I think that intuitive judgments about conscious experience rely not just on one's overall theory, but that

conscious experience itself is a source of evidence, accessible only from the first-person perspective. I think of this evidence as also a kind of intuition, though here more like a Kantian intuition. The idea is that conscious experience itself provides data for theorizing that cannot be derived in any other way. The intuitions themselves - in this Kantian sense - do not automatically yield judgments, so any judgments based on them are going to be susceptible to criticism just like any other inferences based on any data. Nevertheless, theories of conscious experience do need to accommodate this data, or provide compelling reasons for disregarding it.

Now Rosenthal claimed that our intuition that mental qualities are necessarily conscious derives from a theory - maybe a folk theory - about conscious states. The very idea that the first-person perspective provides access to what cannot be otherwise accessed is itself part of the folk theory Rosenthal wants to replace with his theory (as do virtually all materialists). I agree that there is theory - or judgment - involved here. However, what I claim is that it isn't merely a matter of theory, because I claim the first-person perspective, as constituted by conscious experience, provides data. Is that claim itself merely another theory? Well now it becomes difficult to know how to answer that. Here is the data, right in front of me, as it were. Are the inferences I make based on it defeasible? Of course. Is there some sense in which I could be wrong that this data even exists? I suppose, as it is with all other data. (Though I do think it's reasonable to expect that if there is such a thing as first-person data of the sort I'm talking about that it is not susceptible to the same range of defeating considerations that other data are.) But I don't think any of this gives me compelling reason - at least none has been so far furnished - for actually ignoring this data. And to my mind, it supports the view that mental qualities are necessarily conscious. So, that brings me to section 2, the nature of mental qualities.

Mental Qualities

Rosenthal presents his QS theory of mental qualities, on which the very same qualities we are aware of when enjoying conscious experiences are also the qualities instantiated by our perceptual systems when we unconsciously perceive, as in priming, masking, and blindsight experiments. The basic idea is that the external properties we detect - or our perceptual systems detect - constitute "quality spaces" within which specific values along the various dimensions uniquely determine a property within the space. Our perceptual systems detect these values by occupying states that themselves constitute spaces isomorphic to the external property spaces. So our mental qualities are states that pick out specific properties by virtue of occupying the point in the internal space that maps onto the point in the external space of the property in question. To consciously experience a mental quality, on this view, is for us to represent to ourselves that we are in such a state. The state itself, however, is logically independent of any awareness of it. Thus mental qualities are not essentially conscious.

Rosenthal admits that there are intuitive challenges to QS theory, but he dismisses them as mere consequences of a bad folk theory. So, for instance, QS theory is incompatible with the possibility of qualia inversions. Another consequence that seems counterintuitive is that no two creatures - or no creature at two different

times - can share qualitative experiences unless their QS's (for the relevant modality) are the same. If someone is red-green colorblind, for instance, then their experiences of unique blue would not be qualitatively the same as mine.

If either of these scenarios is possible then QS has to be wrong. Rosenthal discusses the inversion case only, but I assume his response to both is the same: it's a "mere intuition", based on folk theory, that these scenarios are possible. QS better accounts for the non-tendentious data - specifically from psychological experiments - so it is to be preferred. What are the non-tendentious data? Basically it amounts to evidence that the same structure of discriminative abilities manifested during conscious perceptual experience - out of which the relevant quality space is constructed - is also manifested during subliminal perception. Since the properties being discriminated and detected during non-conscious perception are the same as the ones discriminated during conscious perception, then it's reasonable to assume that the same mechanisms are involved. In particular, it's reasonable to infer from the similar structures of the two sets of discriminative abilities that they both involve the perceptual system's occupying states within the relevant quality spaces. If so, then it seems as if the very same mental qualities - the same qualitative states - must be involved in both conscious and unconscious perception.

I will grant that we have good reason to believe that in both conscious and non-conscious perception the very same external properties are detected and that the mechanism underlying the relevant discriminative ability involves the same internal states. That our perceptual systems embody a quality space structure that mirrors the structure of the space of the properties discriminated is a view I can accept (assuming it's the best theory on empirical grounds, but let's grant Rosenthal that). But all of this is perfectly consistent with the following hypothesis: When perceptual detection - via the quality space mechanism described by QS - yields conscious experience, a new feature is introduced, the essentially conscious mental quality we experience. This quality - say the redness experienced when looking at a ripe tomato - inherits its relational features from the underlying quality space state, but it adds something else, something essentially experiential. Thus it's not surprising, on this hypothesis, that non-conscious perception should involve the same relational structure as conscious experience, since the relational features of the conscious mental qualities are determined by the structure of the quality space embodied in the perceptual system. Nevertheless there is an intrinsic aspect to these mental qualities that is not captured by the quality space structure, nor is it present when perception is not conscious.

Let's call my view "CE", for "conscious essentialism". CE can explain the same data about unconscious perception as QS, and it is consistent with the possibility of inversions and other sorts of materialistically challenging phenomena. QS explains all the third-person data too, but is inconsistent with the possibility of inversions and the like. On the other hand, CE can seem quite ad hoc. The way it explains the third-person data is by quite shamelessly incorporating QS (or any other competing theory perceptual psychologists finally settle on), and then adding on the essentially conscious features. (As if QS provided a grid for painting by the numbers, and CE takes that and then paints the colors in. On QS, all there are are the numbers and the grid.)

CE does seem ad hoc, and if all we had to go on were the third-person data it would be obviously otiose to posit these extra essentially conscious mental qualities that are parasitic on the QS properties/states. But this is precisely where the appeal to first-person data comes into play. It seems to me that something quite unique and remarkable is presented to us in conscious experience that we encounter nowhere else: what I will call “pure qualities”. By pure qualities I mean properties that are neither structural nor functional (or not only) and have a genuinely substantive and determinate character. Now I have an idea of such a property, and I really can’t see where else I would derive this idea if not from experience itself. What’s more, I seem to really find them there. So, I have a conception of a pure quality, I don’t find any evidence of such a property except from within conscious experience, so maybe this is good reason to think that conscious experience involves features that are present nowhere else. (There is more to say about how this argument goes, but I don’t have the space to fill it out here.)

To sum up. CE accounts for all the same data as QS, respects the inversion and other similar intuitive judgments, accords with first-person data, and indeed explains the inaccessibility of supporting data from the third-person point of view. Does it lead to problems? Of course it does! But I don’t accept that first-person data is ignorable just on the grounds that it makes trouble for theorizing. I claim the first-person data is there - right in front of us, as it were - and it may mean the world is just a very difficult place to figure out.

Higher-Order Misrepresentation

At the end of his discussion, Rosenthal addresses the question of how Higher-Order theories should handle the possibility of internal, higher-order misrepresentation. I just want to say a few words about that issue, as I know there has been a lot of discussion of it in the literature. The first person I know of to have brought up the problem is Karen Neander, and my discussion of it derived from hers. As Rosenthal explains, the problem is how to think about situations in which the higher-order state misrepresents the lower-order one - or, also possible, situations in which there just is no relevant lower-order state. Let’s call these “internal illusions and hallucinations”. The question is: what is it like for the subject when such situations occur?

The answer Rosenthal gives is that what it’s like for the subject is completely a matter of what she is aware of, and therefore is totally dependent on the character of the higher-order representation, as this is the state that embodies her awareness. I think this is clearly the right thing for the higher-order theorist to say. But I do think it undermines some of the initial plausibility of the theory.

Originally, the higher-order view was advertised as a relational view. There is the state one is conscious of and there is the state that constitutes the consciousness of it. With two states and a representation relation holding between them, there were lots of moving parts to the theory that could divide up the explanatory work. One is conscious of the redness of the tomato. The redness is an internal state with whatever features are necessary to constitute a mental quality (on QS theory it’s occupying the appropriate location in the structure, but QS is not required by higher-order theory). The representation relation constitutes the awareness relation, and the subject’s

instantiating the higher-order state constitutes the subject's awareness. Why is consciously experiencing a red tomato so different from merely thinking about one? Well, here one appeals to whatever special features - such as non-conceptual content, etc. - attach to the lower-order state.

But now it turns out that though which lower-order state one occupies will make a difference to a lot of one's discriminative behavior, it doesn't directly (only causally) determine what it's like; that feature depends only on how the higher-order state represents the lower-order one (if it's even there). So it could be like what it's like for me to see the ripe tomato on the kitchen counter even though I'm not in any relevant perceptual state (indeed, suppose my visual cortex has been removed and I'm having "phantom" experiences). I find this a bizarre possibility. But the main point I want to make here is that one thing we wanted from an account of sensory consciousness is a way of drawing a principled distinction between sensorily experiencing something and merely thinking of it. Now, though, it seems to me that we've lost our hold on what could constitute the distinction. If we could appeal to the character of the lower-order states, as in the full-fledged relational version of the theory, then we'd have something to work with (though I still wouldn't be happy, but that's another story). But that feature of the view is gone now.

David Rosenthal has presented us with a rich theory of conscious experience. I take issue with a lot of it, but I want to end by expressing my appreciation of the role his theory has played in deepening our understanding of consciousness.