



The myth of phenomenological overflow

Richard Brown

CUNY LaGuardia, 31-10 Thompson Ave., Long Island City, NY 11101, United States

ARTICLE INFO

Article history:

Available online 22 July 2011

Keywords:

Higher-order thought
Consciousness
Phenomenological overflow
Mesh argument
Methodological puzzle

ABSTRACT

In this paper I examine the dispute between Hakwan Lau, Ned Block, and David Rosenthal over the extent to which empirical results can help us decide between first-order and higher-order theories of consciousness. What emerges from this is an overall argument to the best explanation against the first-order view of consciousness and the dispelling of the mythological notion of phenomenological overflow that comes with it.

© 2011 Elsevier Inc. All rights reserved.

In this paper I will examine the dispute between Hakwan Lau, Ned Block, and David Rosenthal over the extent to which empirical results can help us decide between first-order and higher-order theories of consciousness.¹ According to the higher-order view, having a representation in the relevant first-order brain area (e.g. V1, V2, V3) is not enough for there to be a conscious experience. For that to happen, one needs to also have higher-order representations that are about the first-order representations (presumably in the prefrontal cortex). For instance, on Rosenthal's specific version of the theory (Rosenthal, 2005), in order to make a first-order perceptual state conscious, one needs to have a thought-like state to the effect that one is in that particular perceptual state. This is because to be conscious of something one needs to mentally represent it. Rosenthal calls this the Transitivity Principle which states that "mental states are conscious only if one is in some way conscious of them," (Rosenthal, 2005, p. 3). In contrast a first-order theory maintains that one can have conscious perceptions in the absence of any kind of higher-order awareness. I will argue that Block's mesh argument (Block, 2007, 2008), all things considered, actually supports the higher-order approach to consciousness, rather than the first-order view that Block introduced it to support.

Lau presented data he took to provide evidence for a version of the higher-order theory.² Following standard signal-detection theory (Green & Swets, 1966), Lau distinguishes between d' (pronounced 'd prime') and perceptual certainty. D' is a measure of how much information is captured by a given system. We can think of a person as a signal transducer and see how much information versus noise there is in the signal. Lau presents subjects with stimuli that are just at threshold, so that the visual stimuli are designed to be just barely consciously visible. Participants are then asked to press a button when they have seen the visual stimuli. That enables us to see how successful subjects are, as measured by their d' score. A high d' score indicates that they have a high success rate discriminating the stimulus while a low d' occurs when they have a low accuracy rate. In addition, we can also ask subjects how certain they are that they have seen a stimulus; we can call this second result judgments of perceptual certainty.

E-mail address: onemorebrown@gmail.com

¹ Lau's paper and commentaries on it were originally presented at the NYU workshop on perception October 31st, 2009. This session was video recorded and used with permission. The video of Lau's original keynote and those of the commentaries, and discussion can be viewed at Consciousness Online: <http://consciousnessonline.wordpress.com/2010/02/19/sensory-awareness-and-perceptual-certainty/>.

² On Lau's specific version of higher-order theory (Lau, 2008) the higher-order state consists of an estimate of the probability of being in the first-order state or, as he prefers to put it, an estimate of whether the first-order signal represents something real or is just noise. The differences among various versions of the higher-order theory will not concern us here; our current focus is simply whether there is evidence that higher-order activity of some kind is responsible for consciousness.

Lau's basic idea is that d' gives us a measure of what first-order sensory states there are in the person while perceptual certainty judgments give us a measure of the higher-order state that is present. In a recent study Lau used transcranial magnetic stimulation (TMS) to disrupt functioning in the prefrontal area (Lau & Passingham, 2006), and this had the result of lowering perceptual certainty while leaving d' unaffected. Importantly they were able to find two points with the same d' scores and a difference in the perceptual certainty judgments. That is, subjects were matched with respect to discriminating the stimulus but with TMS stimulation they were relatively unsure whether they had seen anything or not. Without the TMS subjects had higher perceptual certainty judgments. The fMRI data revealed that there was a difference in the activity of dorsal lateral prefrontal cortex (dlPFC) but not in early visual areas in association with the difference in subjective perceptual uncertainty. And on the higher-order hypothesis, visual areas would subserve first-order visual states, whereas dlPFC would subserve a higher-order awareness, in this case related to degree of perceptual certainty. These results seem to suggest that higher-order activity of some sort is necessary for phenomenal consciousness. When the prefrontal area is disrupted, subjects lose their perceptual certainty but are still able to process the stimuli unconsciously. It is plausible to think that the first-order state is responsible for their performance on the task while the higher-order state is responsible for their awareness of the first-order state. Without the relevant higher-order state subjects have no conscious experience.

Block responds that this is not evidence for the higher-order approach. The reason is that he has an alternative interpretation of the data. On Block's view there can be phenomenally conscious states that one does not have cognitive access to. In the case Lau presents, subjects have a phenomenally conscious sensory state, but are unable to access it and so unable to report it. dlPFC is disrupted in Lau's experiment, and since this is the area of the brain we would expect to subserve higher-order judgments about one's experience, that disruption explains why they are unable to report their experience. To many this sounds like a contradiction but Block has mounted a more general argument that there can be phenomenological overflow or that there is more in our conscious phenomenology than we can access. He appeals to the classic Sperling results (Sperling, 1960) as well as more recent results from Landman and Sligte (Landman, Spekreijse, & Lamme, 2003; Sligte, Scholte, & Lamme, 2008; Sligte, Scholte, & Lamme, 2009) as evidence. Let us look at each paradigm in turn.

In the Sperling experiments subjects are presented with a 3×4 array of letters which is flashed very quickly on a computer screen. Subjects report that they see a bunch of letters arranged in a block but are unable to say what most of the letters are. Things change, however, if a tone cuing subjects to report one single row of four letters comes after the display ceases to be visually present. In that case they are able to report the letters in the cued row. Block reasons as follows. We have evidence that subjects have an experience of all of the letters. That is the only way that they would be able to report the letters when they are cued (remember, the cue comes after the stimulus is no longer visually present) and subjects report that they see all of the letters. So the best explanation, according to Block, is that they have phenomenally conscious experience that overflows the higher-order awareness that enables reports of the letters specific identities.

However, this is too quick. The higher-order approach has a ready response to this challenge.³ On Rosenthal's view when one sees the Sperling stimulus one has a first-order mental state that represents all of the letters and their spatial relationships. Because of the conditions of presentation, one's higher-order state represents the first-order state as some letters arranged in a block formation. So the first-order state, which represents all the letters, is conscious but not in respect of all its detail; it is conscious only in respect of its more generic properties. That is because, on the higher-order approach, the higher-order state may make us aware of the first-order state in respect to some of its properties but not all of them.⁴ This explanation of the phenomenon has received some experimental support (see De Gardelle, Sackur, & Kouider, 2009). What this means is that in addition to the properties which figure into my conscious experience there are also non-conscious qualitative properties, which are precisely those properties of the relevant first-order state of which I am unaware.

This response is importantly different from that of Stanislaus Dehaene (Dehaene, Nacchache, Sackur, & Sergent, 2006). On Dehaene's view we suffer from the so-called refrigerator-light fallacy. Often when our mental states are conscious we have a rich phenomenology. So we assume that this rich phenomenology is present even when we are unaware of the relevant mental states; that is analogous to assuming the refrigerator light is on even when the refrigerator door is shut, since it is on whenever we can tell directly whether it is on. On Rosenthal's view there is no illusion. Conscious perceptual experience coincides with what subjects are able to report, though there is often additional perceptual content that is not conscious and which subjects cannot report.

On Rosenthal's view what it is like for these subjects is very much like what it is like to be seeing a group of blurry objects. One consciously sees all of the shapes or letters but does not consciously see all of the specific identities of the letters. Any apparent overflow is explained along the following lines. Subjects see some blurry letters and figure one of two things: Either the letters are actually blurry, or they're well defined and it is distance, or something else, that makes them look blurry. Unlike Dehaene's account the proposal is not that subjects are inferring that there is phenomenology they do not have cognitive access to; they are inferring that there may be something that they are not seeing as well as they could.

³ As David Rosenthal argued in his comments.

⁴ Rosenthal has also deployed the same line of response to Dretske's famous use of change blindness against higher-order theories (Dretske, 2007). In so-called cases of change blindness subjects fail to notice changes that happen right in front of them. Dretske argues that in such cases subjects have a conscious experience of the difference without being conscious of the difference, which looks like a violation of the transitivity principle. Rosenthal has responded (2005, p. 114) that we are conscious of the things in respect of which the two scenes differ, but we are not conscious of those things as the things that make those scenes different. Thus there is no conscious state of which we are not conscious of being in only ones of which we are conscious of some but not all of its properties.

So we seem to have arrived at a stand still. Each is able to give a story that is consistent with the empirical data. On the one hand we have Block's model, which includes phenomenological overflow. On that model we have phenomenology going with the first-order systems and access to those first-order systems going with the higher-order prefrontal areas. On the other hand we have Rosenthal's and Lau's higher-order models, which deny phenomenological overflow. The phenomenology is constituted—and so exhausted—by the higher-order access to the first-order states. Yet each model holds that there are two systems, one with high capacity and one with a relatively lower capacity. Thus each model is consistent with all of the empirical data and each licenses a mesh between psychological and neuroscientific results; the only difference between the models is whether it makes sense to describe what is overflowing reportability as phenomenology, as against non-conscious first-order qualitative states.⁵

Block has recently taken to stressing the Landman and Sligte paradigm as a superior way to make the argument for phenomenological overflow. In the Landman set up subjects are presented with an array of rectangular gradient patches in specific orientations ordered in a circle. They then have a grey interstimulus mask after which another array which either does or does not have one of the rectangle's orientations changed (50% of the time it does, 50% does not). The interestingly new tweak in this set up is that the stimulus is set up so as to evoke what Block calls "a positive afterimage". In essence this is a phenomenally conscious experience of the stimulus that persists for a brief period after the actual stimulus is no longer present. As Block notes, Sligte and Landman have found that this is divided into two phases. In the first of these phases the afterimage is indistinguishable from the experience subjects have when viewing the actual stimulus. After a brief period (roughly between 100 and 300 ms; see Block, 2007, p. 490) the experience becomes ghostly and faded but still distinctly visible according to subjects' reports (this phase can last up to four seconds with training (see Sligte et al., 2008)). Subjects are usually very bad at detecting the change unless there is a cue during the interstimulus period consisting of a line which points to the place where the rectangle that will or will not change would be located. When the cue is presented during the period when the conscious experience is reported to persist subjects are much better at detecting whether there was a change.

So in these cases, Block argues, we have clearer evidence for phenomenological overflow. The subjects are having phenomenally conscious experiences and in the first, no cue, condition they cannot report the orientations of more than 4 of the rectangles. With the cue they are much better. So Block reasons that they have a phenomenally conscious experience that represents the arrangement of rectangles and their specific orientations, since we can report the position of any cued rectangle. However there is more that is phenomenally conscious than what we can report since Landman et al. calculated that subjects could track at most the orientations of four of the rectangles (Block, 2007, p. 488). Subjects say that they see all or most of the rectangles so why should not we conclude that they are right and that there is overflow?

Block acknowledges that one response is to deny that the subjects are right when they say that they really do see all of the rectangles and their specific orientations; perhaps, as before, they suffer from a refrigerator light fallacy. Block gives four reasons against this. These are: (1) that subjects are normal, (2) that they say they are basing their judgments on their visual experience, (3) Block finds it hard to picture what the phenomenology is actually like if one denies phenomenological overflow, and finally (4) that there is also some empirical evidence that the visual system constructs an image, thus corroborating subjects' reports that they are making these judgments based on a conscious visual image. This final consideration is motivated by a variation on the classic experiments that suggests that subjects have access to both size and orientation information.⁶ Let us look at these in reverse order.

There may be evidence for some kind of visual image as the basis for subject's ability to perform but there is no evidence that the image is conscious in respect of all of its details. On Rosenthal's model there can be unconscious qualitative character. This is arguably what happens in standard subliminal vision—e.g., masked priming, and also (what is different) in blindsight. When this qualitative character occurs unconsciously there is nothing that it is like for the subject. Rosenthal has argued extensively for what he calls the homomorphism account of sensory qualities (Rosenthal, 2005), which is precisely meant to be an account of what qualitative character could be specified in a way that does not appeal to consciousness. So the mere fact that subjects have access to information about size as well as orientation may show that there is a visual image that is constructed and being used to generate the comparisons but it cannot show that the image is conscious with respect to these properties, which is what is needed.

What about the third consideration? Block has argued that he does not quite get what the alternative is supposed to be. He distinguishes between general and specific phenomenal content. As he says, both parties agree that there is some general content "letters arranged in a grid," or "rectangles arranged in a circle" but what specific content is there? That is, what exactly is it like for the subject as they look at the rectangles but cannot report their orientations? In particular Block is vexed about what the story would be pre and post cue. If the idea is that there is no specific phenomenology before the cue, but there is afterwards that is very odd and also does not seem to match up with what subjects report (i.e. they do not report any weirdness in their experience). But we have already seen the answer that has been given by the higher-order camp. What it is like is like having the experience of a bunch of indeterminate letters arranged in a grid. There is no (conscious) specific

⁵ In fact, Rosenthal has himself suggested that there may be merely a terminological dispute between himself and Block (Rosenthal, 2002). As he says, "If we agree, however, not to worry about which mental phenomena deserve the honorific title 'consciousness,' it may seem that there is nothing left about which Block and I disagree. We might even agree to apply the term 'conscious', in a special sense, to states that exhibit only thin phenomenality [i.e. non-conscious qualitative character]. Though we are in no way aware of those states, being in them does result in our being conscious of various things. So those states do have an essential connection with consciousness (Page 660)." (Rosenthal, 2002, p. 660)

⁶ See page 309 in Block (2008) for details on the Landman et al. variation.

content in these cases, on the Rosenthal model. Or to put it in another way, the specific content is the content of the first-order states but we are only conscious of them in some respect but not others. So, in answer to Block's question about what it is like for one to be a subject in these experiments on the higher-order approach we can say that it is similar to what it is like when things actually are fuzzy—when one has only some detail because of the optical situation. There is a similar what it is like in the two cases. When the cue comes one's attention is drawn to the location of the rectangle and one is able to make the discrimination.

What should we say about these first two considerations? We do have some reason to take the reports of subjects in the experiments seriously. So we do have some reason to believe them when they say that they are basing their judgments on their conscious visual phenomenology. Subjects also say they have a conscious experience of the shapes but deny that they have conscious experience of the identity of all 12 letters or the orientations of all the shapes. So far nothing in what the subjects report suggests that there is any kind of overflow. In fact at this point it is hard to see how they disagree. Loosely paraphrasing Block says that what is conscious but not reportable are the specific identities of the letters or the specific orientations of all of the rectangles (2008). Rosenthal has agreed with him. One is consciously seeing all of the rectangles, just as one does when one sees a fuzzy image, but fails to be conscious of the specific orientations of the rectangles and so cannot report them. This fits exactly what the subjects report. Block argues that subjects base their reports on a conscious visual image of the stimuli. Rosenthal agrees. They consciously see all of the blocks but fail to be conscious of them in respect of all of their properties (i.e. the specific orientations of each rectangle). Given this we have again failed to find any reason to postulate phenomenological overflow.

Lau appeals to [Rahnev, Maniscalco, Huang, and Lau \(2009\)](#) as providing evidence for the claim that subjects are mistaken when they say that they see each and every rectangle and its specific orientation. In these experiments subjects were asked to attend to an area and were presented stimuli in the unattended region. At the same time subjects were asked to make judgments about perceptual certainty. What they discovered was that when the d' of subjects was matched between the attended and unattended conditions subjects gave a higher rating to their perceptual certainty. That is, they felt more confident that they had seen the stimulus in the unattended region even though they were no better at detecting the stimulus.⁷ This would explain why it is that people feel as though they have more in their conscious experience than they can report. Given this there is no reason to posit phenomenological overflow.

Block responds by arguing that these results actually show that the identification of perceptual certainty with phenomenology is mistaken.⁸ He does not elaborate on the reason, but it should be clear what his reasoning would be. If we identify phenomenology with the activity in the lower visual areas and we find out that those visual areas are incapacitated or otherwise inactive, we should conclude that there is no, or at least very diminished, phenomenology. But it is important to stress that this is on the assumption that phenomenology goes with first-order states, that is, with activity in visual cortical areas. If that were the case then we would have a disassociation between Lau's notion of perceptual certainty and phenomenal consciousness. But this already assumes that Block has successfully argued for phenomenological overflow, and Lau had invoked the Rahnev results to argue against that overflow notion by showing that his model could account for any temptation to posit phenomenological overflow.

We should also note that the same kind of reasoning that Block endorses suggests that we should take the reports of subjects in the Rahnev experiments seriously as well. In Rahnev forthcoming it is noted that subjects are themselves very surprised that they were wrong. This is taken to suggest that they really feel as though they were seeing things in the unattended location even though the first-order activity was very low, suggesting that they did not have the relevant first-order states. Given this it is perfectly understandable why subjects in the Landman and Sligte paradigms say that they see all of the rectangles as well as their specific orientations. Attention is directed to the location of the cue and because of this the Rahnev effect kicks in and subjects overestimate how much they see in the unattended region.

Overall, then, we have seen that there is no reason to believe in phenomenological overflow. It is not entailed by the reports given by subjects in Landman and Sligte experiments nor is it entailed by reports from the Sperling cases. We can explain why people may feel as though they see more than what they can report. Given this the higher-order approach seems to provide the better mesh. It is able to make sense of the putative cases of phenomenological overflow without postulating any such overflow and it is able to explain why subjects report what they do in Rahnev cases. The opposing view has to claim that subjects in the Rahnev cases are wrong when they say that they saw things in the unattended location. This will save one's view but it seems unlikely for the very reasons that Block gives. Given this the overall best explanation would be the one that rejects phenomenological overflow; we simply have no reason to think it exists and no need of it in our theoretical machinery.

As noted at the beginning of this paper, many people find the idea of phenomenal consciousness in the absence of any kind of awareness to be close to contradictory. I have not made much of this in my previous argument since I wanted to evaluate the empirical evidence for the claim. After all, nature retains the right to trump our intuitions. But now that we have seen that there is no reason to postulate phenomenological overflow we should briefly re-examine the issue. The strangeness

⁷ This is what Lau calls Fake Phenomenology, which is roughly the idea that the higher-order state represents the first-order state as being determinate but not representing the details (cf [Grush, 2007](#)). So in the Landman case Lau might hold that the higher-order states represents something like "I am seeing a bunch of rectangles arranged in a circle and I am aware of each rectangle's orientation." In the Landman cases the relevant first-order state is actually present, in the Rahnev cases it is not.

⁸ In his video commentary.

is brought out by asking what it could mean to say that there is phenomenology that we have no cognitive access to? On the face of it, it sounds like we are being asked to believe that there are phenomenally conscious mental states that one is in no way aware of being in but yet there is none the less something that it is like to be in that state. But what sense can we make of a phenomenally conscious state that one is in no way aware of being in? A state is phenomenally conscious when there is something that it is like *for* the creature to be in that state. When there is nothing that it is like for the creature then there is no phenomenal consciousness. So when we are told that there are phenomenally conscious states that we are in no way aware of what we are really being told is that there is something that it is like for me and that I am in no way aware of what it is like for me. This sounds like a contradiction. How can it be like something for me to have a state when, if asked, I would deny that there is anything that it is like for me?

Block has recently acknowledged that there is awareness of some sort involved in every phenomenally conscious experience (Block, 2007, 2008, 2009) and so has come to accept some version of the transitivity principle. But he has resisted the move to a higher-order form of awareness. Instead he has alternated between a deflationary and a self-representational account of the kind of awareness that he thinks we can say is involved in phenomenal consciousness. Following Sosa (Sosa, 2002) he argues that just as I smile my own smiles, I also experience my own experiences just in the having of them. And perhaps that is all there is, Block argues, to the way in which we're aware of our conscious experiences. He has more recently moved from this deflationary reading of awareness to talking about a 'same-order' or self-representational account of the awareness of our own phenomenal consciousness. On this account the phenomenally conscious experience represents itself in some way, but this same-order self-representation is, according to Block, distinct from any kind of higher-order awareness. Let us quickly examine these two proposals.

The terms, 'same-order' and 'self-representational' have also been used by Kriegel (2009), and so one may think that Block has moved to this kind of view. But this would be incorrect. On Kriegel's view both the first-order and higher-order content contribute to the overall phenomenology of one's conscious experience. As Kriegel formulates it, the first-order state contributes the specific content of one's phenomenal consciousness whereas the higher-order awareness is that which makes it a phenomenally conscious state in the first-place. The two aspects are bound into a unitary mental state by some psychologically real process.

But it is arguable that this is just a variant of the higher-order approach. The higher-order awareness here is a distinct state that has to be bound into a composite with the first order-order state. It is also not clear how stable this position is. If the first-order state can occur without the higher-order awareness then we need to know what it will be like for the subject. For instance, consciously seeing red, on this view, would involve a representation, say in V4 being bound to activity in, say, dlPFC (say, via all of these neurons in different places firing in the same frequency (Brown, 2006). The activity in V4 makes it the case that I am consciously seeing red whereas the dlPFC activity makes it the case that I am consciously seeing anything at all. We can then ask what it would be like for someone who just had the activity in V4. If the answer is 'nothing' then how can the first-order state determine what it is like for me when I am conscious of it? On the other hand if it is like seeing red for me then it is not clear how the higher-order awareness can be the factor that determines that it is a phenomenal experience at all.

By contrast, the kind of self-representation that Block is interested in is stipulated to be distinct from the kind of awareness appealed to by Kriegel or Rosenthal. But if the notion of self-representation at work here is not at least something like that of Kriegel's what could it be? Perhaps it is nothing more than the deflationary notion of awareness? Perhaps the basic intuition that Block has is that the notion of awareness has been such that we get it simply by having the relevant experience in the first place.

The problem with either self-representational or deflationary accounts of awareness is that they do not seem to be able to make sense of the idea of this kind of awareness coming and going. A mental state can be conscious at one time and unconscious at another and if we are to take Block's recent proposal seriously that would amount to me having an experience that self-represents and then does not self-represent; but how is this possible if we do not have something like higher-order awareness in mind? The problem is even worse for the deflationary notion. On the deflationary account I am aware of my experiences just in the having of them; so how could I have an experience and yet fail to be aware of it? The result is that qualitative states are always conscious.

Perhaps those like Block do not worry about this apparent mystery. They may see it as the nature of phenomenal consciousness: As with the notion of a quale, there is no such thing as its occurring without being conscious. But this is hardly a congenial position for those that are physicalists. If consciously seeing red, for instance, is just a certain kind of activation or recurrent processing in early visual areas then we certainly have cases where this activation is too low to generate a conscious experience, or so it would seem and in such cases we would have to say that the state self-represents itself at one time and does not at another time. But what could that mean on the present account? So this notion of self-representation is very mysterious. In the absence of some kind of account and without evidence to the contrary what reason do we have for denying that the awareness is some kind higher-order activity? One can dig in one's heels and insist that there is phenomenological overflow in some substantive sense and that these states represent themselves in some mysterious way, and this will save one's theory, but it begins to become clear that this is an *ad hoc* move made to save a theoretical claim. There is no compelling non-question begging reason to think that there is something that both outstrips our ability to report and threatens the higher-order approach to consciousness.

What emerges from this is an overall argument to the best explanation against the first-order view of consciousness and the dispelling of the mythological notion of phenomenological overflow that comes with it. Block argues that if subjects are

right when they say that they see all of the objects then there must be some phenomenally conscious experience that they cannot report. But this is a myth because it relies on the assumption that subjects must have specific phenomenal contents in order for what they say to be true. If Rosenthal is right that the subjects are aware of all of the shapes but not in respect of all of their properties then taking their reports seriously does not entail that there is phenomenological overflow. All that overflows may be non-conscious qualitative character. Secondly any temptation to posit specific phenomenological overflow is mitigated by results like those reported in the Rahnev cases. That data suggests that when we attend to one location we subjectively inflate the phenomenology in the unattended location such that one would feel as though one had seen all of the shapes and their specific orientations even if one had not. Finally, some kind of awareness is required for phenomenal consciousness to account for the distinctive ‘for-me-ness,’ or “me-ishness” as Block has called it. Once one has accepted that this is the case, the most plausible candidate is some kind of higher-order awareness. To put it the other way, what reason is there to prefer the self-representational view? Without a worked out account of what notion of self-representation is in play and how it works, it is hard to say. On the other hand, there is ample empirical and philosophical support for the higher-order approach (Brown & Lau, forthcoming; Lau & Rosenthal, forthcoming) and it elegantly explains the notion of awareness in play in terms of higher-order states representing the relevant first-order states. In short, then, phenomenological overflow is an interesting theoretical possibility but we just do not have any evidence for it and there are better, less mysterious, options on the table.⁹

References

- Block, N. (2007). Consciousness, accessibility, and the mesh between psychology and neuroscience. *Behavioral and Brain Sciences*, 30, 481–548.
- Block, N. (2008). Consciousness and cognitive access. In *Proceedings of the Aristotelian Society* (Vol. cviii, Part 3).
- Block, N. (2009). Comparing the major theories of consciousness. In M. S. Gazzaniga (Ed.), *The cognitive neurosciences IV*. MIT Press.
- Brown, R. (2006). What is a brain state? *Philosophical Psychology*, 19(6), 729–746.
- Brown, R., & Lau, H. (forthcoming). In A. Pautz & B. Stoljar (Eds.), *The emperor's new phenomenology? Empirical support for empty higher-order representations*. Themes from Ned Block. MIT press.
- De Gardelle, V., Sackur, J., & Kouider, S. (2009). Perceptual illusions in brief visual presentations. *Consciousness and Cognition*, 18(3), 569–577.
- Dehaene, S., Nacchache, L., Sackur, J., & Sergent, C. (2006). ‘Conscious, preconscious, and subliminal processing: A testable taxonomy’. *Trends in Cognitive Sciences*, 10, 204–211.
- Dretske, F. (2007). What change blindness teaches about consciousness. *Philosophical Perspectives*, 21(1), 215–222.
- Green, D. M., & Swets, J. A. (1966). *Signal detection theory and psychophysics*. New York: Wiley.
- Grush, R. (2007). A plug for generic phenomenology. *Behavioral and Brain Sciences*, 30, 504–505.
- Kriegel, U. (2009). Self-representationalism and phenomenology. *Philosophical Studies*, 143, 357–381.
- Landman, R., Spekreijse, H., & Lamme, V. A. F. (2003). Large capacity storage of integrated objects before change blindness. *Vision Research*, 43(2), 149–164.
- Lau, H. (2008). A higher order Bayesian decision theory of consciousness. *Progress in Brain Research*, 168, 35–48.
- Lau, H., & Rosenthal, D. (forthcoming). Empirical support for higher-order theories of conscious awareness. *Trends in Cognitive Science*.
- Lau, H., & Passingham, R. (2006). Relative blindsight in normal observers and the neural correlate of visual consciousness. *Proceedings of the National Academy of Sciences*, 103(49), 18763–18768.
- Rahnev, D., Maniscalco, B., Huang, E., & Lau, H. C. (2009). Inattention boosts subjective visibility: Implications for inattentional and change blindness. *Journal of Vision*, 9(8), 157.
- Rosenthal, D. (2002). How many kinds of consciousness? *Consciousness and Cognition*, 11, 653–665.
- Rosenthal, D. (2005). *Consciousness and mind*. New York: Oxford University Press.
- Sligte, I. G., Scholte, H. S., & Lamme, V. A. F. (2008). Are there multiple visual short-term memory stores? *Plos One*, 3(2), 1–9.
- Sligte, I. G., Scholte, H. S., & Lamme, V. A. F. (2009). V4 activity predicts the strength of visual short-term memory representations. *Journal of Neuroscience*, 29(23), 7432–7438.
- Sosa, E. (2002). Privileged access. In Q. Smith & A. Jokic (Eds.), *Consciousness: New philosophical perspectives* (pp. 273–294). Oxford University Press.
- Sperling, G. (1960). The information available in brief visual presentations. *Psychological Monographs*, 74, 1–29.

⁹ Thanks to David Rosenthal, Jacob Berger, Hakwan Lau, Ned Block, Bruce Bridgeman, and Pete Mandik for very helpful comments on previous drafts. Earlier versions of this paper were presented in the Psychology Department at Columbia University March 22, 2010 and at the Association for the Scientific Study of Consciousness June 24th 2010 as “Consciousness and the Tribunal of Experience”.