

On Skepticism about Unconscious Perception

Journal of Consciousness Studies 26(11-12): 8-32(25)

Jacob Berger

Department of Philosophy, Lycoming College

Myrto Mylopoulos

Department of Philosophy and Institute of Cognitive Science, Carleton University

[Please quote the published version.]

Abstract

While there seems to be much evidence that perceptual states can occur without being conscious, some theorists recently express skepticism about unconscious perception. We explore here two kinds of such skepticism: Megan Peters and Hakwan Lau's experimental work regarding the well-known problem of the criterion—which seems to show that many purported instances of unconscious perception go unreported but are weakly conscious—and Ian Phillips' theoretical consideration, which he calls the 'problem of attribution'—the worry that many purported examples of unconscious perception are not perceptual, but rather merely informational and subpersonal. We argue that these concerns do not undermine the evidence for unconscious perception and that this skeptical approach results in a dilemma for the skeptic, who must either deny that there is unconscious mentality generally or explain why perceptual states are unique in the mind such that they cannot occur unconsciously. Both options, we argue, are problematic.

Keywords: perception; consciousness; problem of criterion; masking; perceptual processing

1. Introduction

While there would seem to be much commonsense and experimental evidence that perceptual states can occur without being conscious,¹ such as studies involving so-called 'masked priming' or pathological conditions such as blindsight (see respectively, e.g., Marcel 1983; Weiskrantz 1986; for an overview, see, e.g., Berger 2014), recently some philosophers and cognitive neuropsychologists have expressed skepticism about much of the putative evidence for unconscious perception (see, e.g., Peters & Lau 2015; Phillips 2016a; Phillips & Block 2016; Peters et al 2017; Phillips 2018; Shanks 2017; Knotts et al 2018). Such skepticism comes in roughly two forms: suspicion that purported examples of unconscious perception are instead (a) cases of weakly conscious perception, or (b) cases of unconscious sensory registration that fall short of genuine personal-level perception.

We explore here arguments along both lines. Perhaps the most powerful evidence in support of conclusion (a) comes from a series of experiments by Megan Peters and Hakwan Lau (henceforth

¹ While we do not offer here a complete theory of the nature of perceptual states, we take them minimally to be personal-level states that modally encode information about the body or environment (e.g., Burge 2010).

‘P&L’) (2015) regarding the so-called ‘problem of the criterion’. According to this worry about much experimental work, when participants fail to report seeing stimuli that there is evidence that they did see, it may not be that the stimuli were unconsciously perceived, but that they were consciously perceived but unreported due to high criteria of confidence (for review, see, e.g., Peters, Ro, & Lau 2016). And P&L’s experiments purport to show that standard experimental techniques often thought to render stimuli invisible to consciousness in healthy participants fail to do so for this very reason (cf. Knotts et al 2018). P&L’s considerations, moreover, dovetail with what may be the most well-known challenge in support of conclusion (b): theoretical arguments recently pursued by Ian Phillips (2016a; his contributions in Phillips & Block 2016 and Peters et al 2017; 2018), which revolve around what he calls the ‘problem of attribution’ (e.g., Phillips 2018, p. 481)—that is, the worry that cases of unconscious perceptual processing often thought to be instances of unconscious perception do not meet the requirements for genuine perception.² This pair of challenges may seem to entail either that there is little or no current evidence of unconscious perception or even that there can be no effective study of unconscious perception.³

Whether or not these critiques succeed is hardly inconsequential. Not only does folk psychology at least seem to allow for cases of unconscious perception,⁴ but also perceptual psychology has been professedly studying unconscious perception for decades (e.g., Peirce & Jastrow 1884; Marcel 1983). Similarly, claims of unconscious perception are commonplace in neuroscience (e.g., Lau & Passingham 2007; Seitz et al 2009). If these critiques were successful, we would arguably need to revise both our folk psychological understanding of perception, but also reinterpret huge swaths of experimental results.

Moreover, if it turns out that we have less evidence or even no evidence for unconscious perception, then many *theories of consciousness*, which seek to explain the difference between conscious

² Phillips too discusses the problem of the criterion in various places and regards it as a reason to be skeptical of evidence purporting to establish unconscious perception (see especially Phillips 2016, sections 3 and 4, and Phillips 2018, sections 3.2 and 3.3).

³ See also the recent challenge by Shanks (2017), which raises methodological issues regarding how to assess data purporting to establish the existence of unconscious perception. Shanks’ concern is that selective *post-hoc* data analysis is largely responsible for purported dissociations between performance and awareness. While we think that addressing methodological concerns such as Shanks’ is doubtless important, we focus here on responding to P&L’s and Phillips’ more direct critiques of the evidence for and existence of unconscious perception.

⁴ For instance, ordinary people often say things that reflect a view on which we can perceive aspects of our environment without conscious awareness, such as: “I wasn’t aware of that that loud construction outside until just now, but it explains why I’ve felt distracted this whole time.” There is also some experimental evidence that laypeople believe that “unfelt”—that is, unconscious—pains exist (see Reuter and Sytsma 2018). Though pains may not be perceptual states, these findings lend credence to the view that the folk think of unconscious mentality as relatively widespread.

and unconscious states (e.g., Dehaene et al 2006; Lau & Rosenthal 2011), may be rendered questionable at best. Many of these theories take as their starting point the assumption that unconscious states exist—and then seek to explain the difference between conscious states and their unconscious counterparts. This is the standard *contrastive methodology* of consciousness studies (see, e.g., Baars 1997, p. 294). If these skeptical arguments succeed, this method, along with many scientific and philosophical theories of consciousness, would seem to be undermined. Indeed, it may become unclear how to study consciousness at all.

Some theorists have attempted to reply to these critiques in various ways (e.g., Block 2016; Block's contribution in Phillips & Block 2016; Block's and Kentridge's contributions in Peters et al 2017; Quilty-Dunn 2018; D'Aloisio-Montilla 2018), mostly by offering ostensible examples of genuine unconscious perception. Whether or not such cases are compelling, however, we pursue here different lines of response to P&L and Phillips: we argue that their considerations simply do not decisively undermine the existing evidence for unconscious perception. To be clear, we do not demonstrate the opposite: that there is clear evidence of unconscious perception. Our goal here is more modest: to provide a first line of defense against these recent skeptical approaches to unconscious perception by both showing how they fall short in specific ways and by raising some general worries about their implications.

After briefly characterizing how we understand (un)consciousness in section 2, we describe in section 3 what may seem to be a dilemma for the study of unconscious perception generated by P&L's and Phillips' considerations. We then argue in sections 4 and 5 that there is room to resist both horns of this purported dilemma—and that P&L and Phillips do not demonstrate that there is no evidence of unconscious perception. In section 6, we close by arguing that if, contrary to what we argue, the conclusion of the dilemma that we address is accepted, it actually gives rise to a further dilemma, but for the skeptic about unconscious perception. This puts pressure on the skeptic to explain why the initial dilemma should not be rejected.

2. On (un)consciousness

To assess critiques of the study of unconscious perception, we naturally must first have a handle on what it is for a perceptual state to be conscious or unconscious. There is, however, considerable debate about how to characterize consciousness, not only within the philosophical literature (see, e.g., Berger 2014), but also within the psychological literature about how to operationalize it (see, e.g., Persuh 2018).

Many philosophers today follow Thomas Nagel (1974) in holding that a mental state is conscious just in case *there is something that it is like* to be in that state. Others—most commonly so-called ‘higher-order theorists’ of consciousness—characterize conscious states as those states of which one is suitably aware of being in (e.g., Armstrong 1968; Lau & Rosenthal 2011). In his work, Phillips explicitly assumes the former gloss (e.g., Phillips 2018, p. 481), though Lau is sympathetic to the latter (e.g., Lau & Rosenthal 2011). We will not adjudicate between these two ways of characterizing consciousness; our arguments apply regardless of which one is adopted. Going forward, we understand an unconscious state to be a state either that there is nothing that it is like for one to be in or that one is not suitably aware of being in.

More important for present purposes is the question of how conscious perception is operationalized within cognitive science. And, again, there is much debate about this. To illustrate how consciousness is typically studied experimentally, consider standard experiments involving the technique of visual masking, which is widely used to study unconscious vision in healthy individuals (for an overview, see, e.g., Bachmann & Francis 2013). In a typical study, stimuli that would otherwise be consciously seen are presented and *masked*—that is, either preceded or followed (or both) by different spatially or temporally congruent stimuli—in ways which ostensibly render the initial targets invisible to visual consciousness.

Why are such target stimuli thought not to be consciously seen? In the past, much work on unconscious perception has operationalized conscious visual states as those states that are assessable by *subjective measures*, wherein participants in some way report that they did or did not see the target stimuli. Such measures are subjective insofar as they require participants to make judgments about their own perceptual states. There are various kinds of subjective measures; participants might, for example, press a button to indicate that the target was perceived or rate how visible the target was on a Likert scale (e.g., ‘0’ for totally invisible and ‘7’ for totally visible).

It would make sense that subjective measures have been typically considered the gold standard measure for consciousness. After all, subjective reports are thought to be typically, though of course not always, reliable indications that one has consciously seen something (see, e.g., Bayne et al 2014, p. 481).

Many theorists have urged that, as indicators of consciousness, subjective measures are preferable to *objective measures*, wherein participants perform a kind of behavioral operation involving targets. Again, there are a variety of such measures. Participants might, for instance, be *above chance* at discriminating the target: for example, they might be able to accurately select the stimuli in forced-

choice scenarios (e.g., they are more likely to select a square from a choice of either a square or circle if first presented with a masked square). In signal-detection theoretic terms, participants' performance on the objective task would be such that their d' , the signal-to-noise ratio, is greater than zero (see, e.g., Stanislaw & Todorov 1999). Some have worried, however, that using such objective measures of task performance as evidence that targets are consciously perceived is problematic since there is the possibility that the objective performance may be driven by unconscious perception (for review, see, e.g., Persuh 2018; Rosenthal forthcoming).

This is why it is often thought that a promising way to demonstrate unconscious perception—in for example masking studies—is to establish a simple *dissociation* between subjective and objective measures. For example, even if participants do not indicate or even deny that they saw target stimuli (subjective measure), they are nonetheless above chance at discriminating the targets (objective measure). Alternatively, even if participants neither report that they saw the target stimuli nor are able to discriminate them—that is, their discriminatory $d'=0$ —they might still be *primed* by the targets insofar as their reaction times are facilitated in subsequent target-related tasks (e.g., they are faster at identifying a square as a square if first primed by a masked square than a masked circle) (see Merikle & Reingold 1988). This kind of procedure for studying unconscious perception has, however, faced much recent criticism.

3. A (seeming) dilemma for studies of unconscious perception

Many theorists have recently come to doubt that performances on standard subjective measures such as ratings on visibility scales are adequate evidence that stimuli have not been consciously perceived. It is well known that experiments that rely only on such subjective measures of conscious invisibility face the problem of the criterion. That is, participants may be conservative about what they are willing to report as having seen, and so fail to report seeing stimuli that are weakly, but consciously, perceived, as evidenced by success on an objective task. Many have questioned studies of unconscious perception because they fail to control for such a possibility (for just one example, see, e.g., Stein et al 2016 on Soto et al 2011).

The confound of criterion bias is only *potential*, however, insofar as it at least seems to remain open that such above-chance discrimination in the absence of subjective report is driven by genuinely unconscious perception. But this leads us to P&L's (2015) experimental work, the results of which may suggest that the problem of the criterion is not merely a possibility, but an actual obstacle to using this methodology to study unconscious perception.

To attempt a demonstration of unconscious perception that avoids the problem of the criterion, P&L employed a two-interval forced-choice procedure with confidence ratings. Participants were presented with two masked displays in subsequent intervals, one involving a target stimulus and one that was, unbeknownst to them, blank. In addition to a discrimination task, in which they were asked to indicate the orientation of the target in both the stimulus-present and blank intervals, participants were asked to bet on the discrimination choice in which they were more confident.

P&L sought to demonstrate an instance of what they call ‘Performance without Awareness’—that is, cases wherein participants were above chance at discriminating stimuli, though equally confident in their discrimination choices in the stimulus-present and blank intervals. P&L’s reasoning is that such performance differences without differences in confidence would indicate that participants genuinely did not consciously perceive the stimuli. As they put it, “observers are unconscious of the information contributing to their decision if they can discriminate a target above chance, but doing so feels no different introspectively from discriminating (or guessing about) nothing at all” (p. 7; see also Peters in Peters et al 2017, p. 2). P&L take any differences in confidence judgments involving the stimulus-present and blank intervals as evidence that the intervals “felt” different—that is, that the stimuli were consciously perceived. In other words, P&L seem to regard, as others recently have, the kind of metacognitive tasks that their study employed as a kind of subjective measure of consciousness (see also, e.g., Persaud et al 2011; for discussion, see Rosenthal forthcoming). And since these confidence ratings can be elicited whether or not participants indicate via a standard subjective measure that they saw those stimuli, a demonstration of Performance without Awareness in P&L’s study would seem to constitute a case of genuine unconscious perception, uncontaminated by criterion bias.

But P&L were not able to demonstrate Performance without Awareness in these conditions. Instead, they found that participants were significantly more confident in their judgments regarding stimuli than in their judgments of blanks when they were able to successfully discriminate the orientation of the target. P&L thus claim that, when participants’ discriminatory d' was greater than 0, they *did* consciously perceive the stimuli to some extent. In other words, they suggest that, if we control for criterion bias, we should expect no dissociation between subjects’ confidence in their

judgements pertaining to the stimuli and blank intervals using these masking techniques that supposedly render stimuli invisible to visual consciousness.⁵

It is important to clarify that P&L do not deny that unconscious perception exists. For one thing, they (2015, p. 10) claim that blindsight, which occurs in people with severe damage to their visual cortices (e.g., Weiskrantz 1986), constitutes an actual example of Performance without Awareness. One might worry, however, that blindsight is not an uncontroversial illustration of unconscious perception, as some—including Phillips (2016, p. 434-441)—speculate that it may simply involve a form of degraded conscious vision (see also, e.g., Overgaard et al 2008; though see, e.g., Weiskrantz 2009). Indeed, since blindsight is a pathological condition, one might doubt that it can reveal anything about how healthy vision works—and in particular whether healthy vision can occur unconsciously.

Could there be less contentious experimental demonstrations of Performance without Awareness? Perhaps. As P&L (2015, p. 10) themselves note, different experimental techniques, such as those involving the procedure known as ‘transcranial magnetic stimulation’ to induce blindsight-like effects in healthy participants, might mask stimuli in a way that enables successful discrimination without differences in metacognition regarding that discrimination. We note, however, that J. D. Knotts, Lau, and Peters (2018) uncovered similar findings as their (2015) study using the masking technique of continuous-flash suppression. Of course, P&L further observe that their “findings cannot rule out all forms of unconscious perception, such as subliminal priming, in which the evidence for unconscious processing is typically indirect benefits in reaction times” (2015, p. 10). Thus one might still think that masked-priming studies could provide evidence for unconscious perception. But, as P&L also emphasize, their findings do bear on how such studies must be conducted: to avoid criterion bias in these cases, it must be demonstrated that there is genuinely no conscious perception via the use of additional measures such as confidence ratings.

Thus it is hard to see what such an experiment would fully look like, let alone to anticipate whether or not it would reveal unconscious perception. This is in part why further experimental

⁵ We note that some may disagree here, arguing that these confidence judgements are better viewed as objective measures, since they indicate that participants could *detect* the stimuli. On this view, we might say that P&L found that when participants’ *discriminatory* $d' > 0$, so too was their *detection* $d' > 0$, as expressed in their confidence judgments—and that a successful demonstration of Performance without Awareness would thus require detection to be at chance, but discrimination to be above chance. (Although P&L do not use this language of ‘detection’ versus ‘discrimination’ d' , we thank an anonymous referee for another journal for introducing it to clarify the discussion.) Here we simply follow P&L in taking these judgements to be a form of (non-standard) subjective measure, and note that, on their view, unconscious perception would be demonstrated were there a dissociation in terms of such judgements and (discriminatory) d' .

work is necessary to settle the matter. But we are not experimentalists—so we will not pursue this issue further. Our point at this juncture is that one might conclude that there is presently no clear experimental evidence of Performance without Awareness (in healthy individuals) due to the problem of the criterion.

Given all this, one might instead think that the solution is to explore work wherein not only such metacognitive judgments, which indicate that target stimuli have been detected, but also discrimination are at chance. A common experimental paradigm in which these conditions often obtain is that of masked priming, wherein a masked target primes a response to a subsequently presented congruent stimulus despite at chance discrimination and detection of the target. One might think, then, that such studies are the place to look for the existence of unconscious perception (but see, e.g., Lau 2009, p. 163-164; Peters' contribution in Peters et al 2017, pp. 1-2). But this takes us directly to Phillips' challenge. What Phillips argues is that, at that point, no longer can we safely assume that such experiments correctly attribute personal-level perception to participants. This is what he calls 'the problem of attribution'.

Phillips maintains, plausibly, that to demonstrate that there is unconscious perception, we must first understand theoretically what is *genuine perception*. All parties to this debate grant that there are unconscious states involved in the perceptual hierarchy that carry information but that are, to use Daniel Dennett's (1969) expression, *subpersonal*—that is, attributable only to the person's subsystems and not to the individual herself. Retinal states are plainly causally involved in generating visual states, but they are not themselves psychological states. What is at issue, then, is whether or not there are any *personal-level* perceptual states that are unconscious.

What is it for a perceptual state to be personal-level? Like most notions in this area, there is much debate (see, e.g., Drayson 2012). One strategy is to identify some specific functional role that all and only personal-level perception is capable of satisfying. Thus Tyler Burge (2010) notably argues that genuine perception exhibits not only perceptual constancy, wherein distal stimuli are stably discriminated amid variation in environmental conditions and proximal stimulation (see, e.g., chapter 9), but also that it "occurs and figures directly in guiding action" (p. 375). Phillips, following Burge, endorses these requirements too.

The problem for many studies of unconscious perception, Phillips argues, is that the perceptual states involved do not seem to play such roles. In the masked-priming studies that are relevant for Phillips' purposes, subjective and objective measures are congruent—that is, participants report that target stimuli are invisible and discrimination is at chance—and yet participants are

primed by those stimuli. But since participants cannot discriminate the targets in such studies, Phillips urges that the subsequent priming effects could be explained without appeal to personal-level perceptual states. It could be, for example, that participants are faster at identifying a square as a square if first primed by a masked square than a masked circle because their visuo-motor systems come to carry information about squares in virtue of the prime, not because they genuinely perceive a square. In other words, we might question whether such studies demonstrate unconscious perception because, as Phillips puts it, such states seem “confined to the visual system and so wholly unavailable for action control and guidance” (Peters et al 2017, p. 5). Phillips therefore concludes that much study of unconscious perception may not in fact investigate unconscious perception at all, but rather mere subpersonal sensory processing.

We are now in a position to see how P&L’s and Phillips’ combined considerations seem to give rise to the following skeptical dilemma. Here is how Phillips recently summarizes the concern:

We arrive then at a stark dilemma for the proponent of unconscious perception. For it is very hard to see how any study could avoid both the problem of the criterion and the problem of attribution. We can sharpen the dilemma by noting that above chance discriminative responding very plausibly operationalizes a basic requirement for individual attribution. If that is right, then effects at the objective threshold cannot possibly provide evidence of perception proper, and the problem of the criterion is unavoidable (2018, p. 26).

In other words, due to the problem of the criterion, P&L’s findings might lead us to expect that there will be no experimental evidence (in healthy individuals) of discrimination above chance without conscious perception. One might thereby think that to demonstrate unconscious perception, studies must use objective measures of conscious invisibility wherein discriminatory $d' = 0$, as in certain cases of masked priming. The difficulty then is that such studies face Phillips’ problem of attribution: that the effects found in these conditions may be due to subpersonal processing, not genuine perception.

Phillips characterizes this dilemma as a problem for the proponent of unconscious perception, insofar as it seems to entail that there can be no meaningful evidence of unconscious perception. We argue here, however, that (i) there are reasons to think that this dilemma can be successfully addressed, and (ii) in any case, the dilemma should rather be conceived of as a challenge for the skeptic about unconscious perception, not its proponent.

4. The first horn: the problem of the criterion

Whether or not experiments that conclusively demonstrate Perception without Awareness are possible, there are good reasons to doubt that P&L's experimental work establishes their conclusion.

Recall that P&L base the logic of their interpretation of their results on the core assumption that if an observer's subjective experience of a stimulus is different from the subjective experience of nothing at all, then it is consciously seen. But notice that this does not entail that one's subjective experience of the stimulus is different due to one's being in a conscious *visual* state. As P&L themselves acknowledge, one way to respond to their study is to maintain that participants' reports may not be due to conscious perception, but to "some sort of non-visual 'hunch' or 'feeling'" (2015, p. 7). That is, it may be that participants' unconscious perception of the stimuli causes them to have nonvisual states—hunches—that register that a stimulus was present, which in turn causes participants to report that they are more confident in those trials.

This alternative hypothesis has an obvious connection with blindsight. In cases of so-called "Type-2" blindsight, individuals deny consciously seeing stimuli presented in their blind field, but are nonetheless able to make above-chance "guesses" as to, for example, the shape, location, and color of the presented objects, based on what they acknowledge to be some form of non-visual awareness of the stimulus. It is reasonable to construe these "guesses" along the lines of the non-visual hunches we are positing to explain the results of P&L's study. The general idea here is that nonconscious perceptual states can sometimes be monitored or tracked by non-visual states that are robust enough to inform metacognitive judgements, but not verbal reports pertaining to the direct perception of the stimuli. If this hypothesis regarding hunches is correct, then we would expect that, even controlling for criterion bias, participants would not report experiences of masked stimuli, since such perceptual states would be genuinely unconscious. Furthermore, the confidence judgments that participants make would not be subjective measures of visual consciousness, as P&L take them to be. Rather, on our view, such judgements would be subjective measures of one's awareness of one's confidence (see also Rosenthal forthcoming, section 3), which would itself be based on a non-visual hunch that one saw something.

P&L object that "this issue is essentially one of terminology" and that "strictly speaking, a non-visual hunch is also defined as conscious so long as it meaningfully tracks visual processes" (2015, p. 7). In other words, one might think that all there is to a visual experience of a stimulus is the conscious impression that some stimulus is/was visually present—what else could such a hunch be if not the experience itself (or a part of it)?

But whether or not such hunches are conscious, we must be careful to distinguish (conscious) *visual* states from (conscious) *nonvisual* states that are *about* visual stimuli. That is to say, there is at least a conceptual difference between the conscious/nonconscious distinction and the visual/nonvisual distinction. An ordinary occurrent thought about what one currently sees (e.g., the thought that there is a square present) may be caused by visual processing, but folk psychology seems not to regard such states as visual experiences themselves. After all, one can have such a thought with one's eyes closed—and arguably without any mental imagery or visualization at all. What might the distinction between visual and non-visual states come to? We plainly cannot defend a theory of the difference between perception and cognition here, but it would seem open that one could slot in any of the many existing theories of that distinction to account for the differences between visual (perceptual) states and nonvisual (cognitive) hunches (for an overview of some theories, see, e.g., Quilty-Dunn 2016). Perhaps nonvisual states represent states of affairs in discursive or conceptual/linguistic formats, whereas visual states present visual stimuli in iconic or nonconceptual ways. On this view, perhaps a non-visual hunch has a conceptual content such as “there is a square,” whereas one's perceptual state has a picture-like content that presents the square in a distinctively visual and nonconceptual way (e.g., that involves color boundaries).

Whatever the final analysis, it is far from obvious that these distinctions track the conscious/unconscious distinction. Nothing about a state's having conceptual or nonconceptual content entails that the state must be conscious—or at least, much more needs to be said to establish such a conclusion. For present purposes, we need not claim that nonvisual states can and do occur (un)consciously—our point is simply that the mere conceptual difference between a state's being visual and its being conscious is sufficient to show P&L's conclusion does not immediately follow.

This is because we can in turn distinguish (conscious) states that track one's unconscious visual states and those that track conscious visual states. It is thereby open that an *unconscious* perceptual state might cause a nonvisual conscious feeling that one saw something, which could account for participants' betting behavior.⁶ Indeed, this non-visual hunch might be reasonably construed as a *metacognitive feeling*—and there is much independent experimental evidence that such feelings can track unconscious states. So-called ‘feelings of knowing’, for example, monitor ongoing

⁶ This alternative hypothesis is even more credible when we recognize that it is likely that such hunches are themselves typically unconscious—and only made conscious upon being asked to render a judgment of confidence. We thank Kate Pendoley for pointing this out to us.

memory retrieval processes without the states involved in those processes being conscious (e.g., Koriat 2000). Similarly, participants in P&L's experiment may have a hunch or feeling that something is present, without being aware of the visual state that gives rise to it. In sum, that P&L found *some* introspective difference between the intervals in their experiment is not evidence that the difference is *perceptual*—and this is the key claim that would need to be established in support of their conclusion.

P&L did run a control study, which may seem to address this alternative hypothesis involving hunches. In these trials, “the subjective task was to indicate which interval appeared more *visible* rather than confidence in the corresponding discrimination” and P&L found that “as soon as participants were able to discriminate the target above chance, they were able to indicate which interval contained the target above chance” (2015, pp. 7-8, emphasis theirs). That is, it would seem that participants regarded the states on which they based their confidence judgments as visual.

But even these trials do not rule out the possibility that the relevant perceptual states were unconscious and that the relevant cognitive impressions were nonvisual. After all, it is not implausible that participants would quickly or automatically infer that the nonvisual impression that some stimuli were present was caused by their having seen those stimuli. Since it is not often the case that we have the impression that we saw things without having consciously seen them, the fact that this experimental set up dissociates these commonly co-occurring mental states may confuse participants to some extent. This alone could account for their judgments that the masked stimuli were more visible. Moreover, since ordinary participants are typically not savvy regarding fine distinctions between types of mental state, it is plausible that they simply mischaracterized those hunches, which do *represent* visual states, as being themselves visual. In other words, participants may simply have misconstrued such hunches, which have *content pertaining to visual states*, as visual states.

One might reply that, on balance, P&L's explanation is nonetheless to be preferred because our alternative explanation requires attributing to participants this sort of confusion. But P&L's experiment puts participants in a highly atypical perceptual situation, and requires them to perform an equally unusual task—and so it should be expected that their awareness of what psychological states they are in might be prone to some error. In addition, an appeal to hunches helps to explain why participants are typically reluctant to report that they saw something in such cases, and are seemingly at best able to make forced confidence judgments regarding their discriminations. In any case, additional study is needed to tease apart these competing hypotheses, which we leave to experimentalists. But it is safe to conclude that P&L offer no dispositive reason to doubt that studies

wherein $d' > 0$, but participants report not perceiving the target stimuli, involve unconscious perception. Controlling for criterion bias, it remains at least open that there is dissociation between ordinary subjective measures of visibility and objective task sensitivity, even if participants render metacognitive judgments of the sort that P&L observed. What of studies wherein discrimination is at chance? Might those involve only subpersonal processing, as Phillips argues?

5. The second horn: characterizing perception

Recall that, according to Phillips, genuine perception must play an action-guiding role, and for this reason, masked-priming results do not point to unconscious perceptual states, as against subpersonal processing. We start by noting that some may view this appeal to an action-guiding role as an overly stringent requirement on genuine personal-level perception. Perhaps a role in action guidance is not strictly necessary, but merely one of many possible *indicators* that a perceptual state is personal-level (for such a proposal, see Block 2016 and Block's contribution to Phillips and Block 2016). But to settle the question of how to draw the distinction between the personal and sub-personal level would take us too far afield here. Instead, we grant this criterion to Phillips.

Yet even if we do grant this conception of perception, and even if there is no masked-priming-based evidence that unconscious states play action-related roles, our first reply is to observe that this lack of evidence is arguably an artifact of our current methods for masking stimuli (e.g., Lau 2009; Persuh et al 2016). As is well known, most of our present techniques for masking stimuli involve degrading their signal strength (e.g., presenting them for short durations). So even if the relevant states involved in masked priming do not drive action, it could be due to the fact that these states encode information weakly, not because they are not genuinely perceptual. Perhaps experimentalists will devise or more fully explore new masking techniques to generate uncontroversial evidence of unconscious perception (for efforts towards one such technique, see, e.g., Persuh et al 2016).

Indeed, when signal strength is not an issue—that is, in non-masking cases—there does appear to be evidence that unconscious perception can control and guide action in the ways required by Phillips' conditions. In his own reply to Phillips, Ned Block (Peters et al 2017, p. 8) offers the example of Mel Goodale and David Milner's (2008) well-known work on dorsal-stream perception in individuals with visual-form agnosia. Persons suffering from this disorder fail in their manual and verbal reports of size constancy, but are nonetheless capable of successfully guiding via vision their

grip aperture around objects, suggesting that they do visually perceive their size, albeit not consciously.

Phillips objects, however, that Milner and Goodale themselves maintain that this so-called ‘vision-for-action’ is, despite appearances, really “not perceptual in nature” (2008, p. 776; quoted in Peters et al 2017, p. 6). So this fails to constitute a genuine case of unconscious perception. Importantly, however, Milner and Goodale are clear that they assume a conception of perception wherein a state is perceptual only if it is either conscious or “potentially *could* reach conscious awareness, e.g. with slightly different stimulus parameters...” (2008, p. 775, emphasis theirs).⁷ Since dorsal-stream states are typically viewed as being inaccessible to consciousness, it would seem that they are not perceptual on Milner and Goodale’s conception.

But an appeal to consciousness in characterizing such states as perceptual is problematic for two reasons. First, since what is at issue here is whether or not perceptual states can occur unconsciously, it is question begging to make it a necessary condition that perception enjoys some relation to consciousness. Second, characterizations of the mental that build in such a relation face independent worries. There are a host of contingent reasons why certain kinds of mental state may not ever reach consciousness, which have little or nothing to do with the functional or causal roles by which we individuate them (see, e.g., Rosenthal 1990). So long as Phillips is individuating perception functionally, it is far from clear that he should endorse Milner and Goodale’s denial that dorsal-stream visual processing is perceptual, especially if such states seem to exhibit other features of perception, such as supporting perceptual discrimination and action-guidance.

Phillips nonetheless offers a seemingly independent reason to regard such vision-for-action as subpersonal—namely, that the type of action control it supports appears more akin to the work of an “automatic pilot” (Peters et al 2017, pp. 6-7) than that which is genuinely attributable to the individual. But the fact that some behavior is automatic does not entail that it is driven by subpersonal mechanisms. As Ellen Fridland (2017) persuasively argues, automatic behavior, as exemplified in skilled action, is not ballistic or invariant, and is coherently integrated with one’s personal-level intentions in a ways that involve sensitivity to the contents of those states. The motor system will, for example, automatically adjust its outputs to compensate for task-relevant but not task-irrelevant perturbations during a task (e.g., Liu & Todorov 2007), where task-relevance here is determined by the content of the agent’s intention. In light of this kind of integration between

⁷ This point is later acknowledged by Phillips (2018, p. 501, fn. 48).

automatic states and processes and personal-level states, such as intentions, there is no good reason to hold that because dorsal-stream vision proceeds automatically, it is thereby subpersonal.

Phillips denies, however, that vision-for-action is sensitive to intention in the relevant ways. He approvingly cites Danckert and Rosetti (2005), who maintain that dorsal stream vision “often functions automatically, rapidly modifying visually guided hand movements... in contradiction to conscious commands” (p. 1042; see also Pisella et al 2000). But Pisella and colleagues’ (2000) study, to which Danckert and Rosetti refer in support of this point, shows only that movements with a time course of under 300 ms are encapsulated from intention. This is consistent with some cases of vision-for-action involving slower movements’ being sensitive to conscious intention, as the same study found. In addition, it is worth noting that the kinds of reaching tasks used for testing individuals with visual form agnosia—our candidate case of unconscious vision-for-action—do not involve such short time scales. So we have no reason to expect the kind of insensitivity to intention exhibited in Pisella and colleagues’ study, and thus no reason to accept that such vision fails to guide action in the way that constitutes genuine perception.

We conclude, then, that Phillips offers no good reason to doubt that the visual states implicated in cases of vision-for-action are genuine, personal-level perceptual states. In the next section, we argue that the seeming dilemma we have been addressing for evidence about unconscious perception results in a challenge for the skeptic about unconscious perception.

6. A dilemma for skepticism about unconscious perception

To take stock, so far the form of skepticism we have been discussing concerns the plausibility of the available evidence for unconscious perception. Call this ‘weak skepticism’. But notice that we can distinguish this kind of skepticism from a more radical form of skepticism about the existence of unconscious perception itself, which we will call ‘strong skepticism’. While weak skepticism clearly does not entail strong skepticism, we argue here that if one accepts the weak skeptical dilemma above, then it would seem that one is forced into endorsing strong skepticism (or something close to it). After all, if one doubts that there is any evidence for the existence of some phenomenon X, then one should be skeptical or at least agnostic about the existence of X. But strong skepticism itself faces a dilemma.

In short, since there would seem to be much commonsense and experimental evidence that *other* kinds of nonperceptual mental states—such as beliefs, desires, and emotions—can occur without being conscious, the strong skeptic must either explain why perceptual states are unique in

the mind insofar as they cannot occur unconsciously or maintain that these other kinds of mental states cannot, despite appearances, occur unconsciously. Both horns lead to troubles for this skeptic.

Consider first that many theorists take there to be solid and growing evidence for nonperceptual unconscious mentality (for an overview, see, e.g., Berger 2014). Folk psychology maintains that people typically have many unconscious beliefs, desires, and emotional states of which they are unaware. We often explain, for example, a friend's erratic behavior regarding a romantic partner as a result of unconscious (though perhaps irrational) beliefs in one's inadequacy, desires for freedom, jealousies, and the like. We also commonly appeal to unconscious states in explaining everyday cases of self-deception, distracted deliberation, and confabulation.

There is much experimental evidence for such states as well. In the case of cognition, for example, the supposition that there is significant mental processing that occurs nonconsciously underlies the commonly endorsed dual-process theories of cognition (e.g., Tversky & Kahneman 1974; Evans and Stanovich 2013), usually proposed as models of reasoning, decision-making, and judgement (see also, e.g., Levy 2011 for application of such models to cases of weakness of will and self-control). According to such views, there are two main types of processing that underlie our decision-making and judgements. Type-2 reasoning is thought to be deliberate, slow, parallel, and typically conscious. We use this kind of reasoning when, for example, we are trying to work through a difficult problem or plan the stages of our day. Type-1 processing, on the other hand, is thought to be automatic, fast, parallel, and operate largely outside of conscious awareness. It is typically associated with heuristic reasoning and even with various cognitive biases (e.g., the conjunction fallacy). Important for our purposes is to emphasize that the steps involved in cases of Type-1 reasoning do not seem to be available to conscious awareness. One may be aware of the information one is reasoning *about*, and be aware that one is engaging in *some sort of* goal-directed reasoning in response to a question pertaining to that information (e.g., "A bat and a ball together cost \$1.05. The bat costs \$1 more than the ball. *How much does the ball cost?*"), but one is not aware of all the states and processes involved in the reasoning itself. This would seem to be a clear example of unconscious nonperceptual mentality.

Likewise, there is well-known evidence that people harbor so-called 'implicit biases', which are putatively unconscious attitudes that influence individuals' social behavior (for an overview, see, e.g., Mandelbaum 2016). Implicit-association tests, for example, reveal that participants who claim not to have racial biases are typically slower at categorizing positive terms when presented with black faces than with white faces (e.g., Greenwald et al 1998). Such results are often explained by positing

that these participants have negative attitudes, of which they are not aware, toward black people that biases their response times. And, as Miao Qian and colleagues (2017) report, implicit biases have been appealed to in explaining “far-reaching negative consequences in all spheres of human life including education, healthcare, employment, justice, finance, dating, and politics” (p. 845).

We acknowledge that there remains much debate about the psychological nature of implicit bias. Some theorists, for instance, maintain that implicit attitudes are ordinary propositional thoughts that are simply unconscious (e.g., Mandelbaum 2016), whereas others propose that they are mere conceptual associations without full-blown propositional structure (e.g., Gawronski & Bodenhausen 2011). While we cannot adjudicate here such debates, our point is that there would seem to be much evidence that there is a kind of cognition that occurs without being conscious. Of course, one might insist that such cognition is not on par with perceptual states for various reasons. For example, it may seem that implicit attitudes are merely dispositional states, whereas perceptions are occurrent states, leaving open the possibility that implicit bias manifests in conscious judgements, while not being conscious itself.⁸ But there is reason to prefer an occurrent reading of the states implicated in implicit bias over a dispositional reading, since they can and often do play action-guiding roles.

If the strong skeptic accepts that it is reasonable to posit unconscious mental states in the cases we have offered, then they must explain why these types of mental state, but not perceptual states, can be tokened outside of consciousness. Offhand, one would expect parity here. Why should one class of mental states always be conscious while others fail to be so regularly? While there are of course some differences between perception and other varieties of mentality, the case for parity is especially strong when one considers the possibility that some form of *intentionalism* may be true, according to which perceptual states are simply intentional states with particular sorts of contents or formats (e.g., Byrne 2001).⁹ If perceptual states and beliefs differ only in, for example, the nature or format of how they represent things, and if beliefs can occur unconsciously, then it would stand to reason that perceptual states can occur unconsciously as well. There is no good reason to think that a difference solely in content or format is relevant to consciousness.

⁸ We thank an anonymous reviewer for raising this point.

⁹ We note that Phillips himself has expressed sympathy for relational or naïve-realist views of perception (see, e.g., Phillips 2014), on which perception need not have representational content, but is instead understood as a direct relation between perceivers and perceived objects. But even on such a view, there is no clear independent reason why the relational nature of perception would be a relevant difference between perception and belief such that the latter but not the former can occur unconsciously.

Similarly, most theories of consciousness currently available provide no good grounds for resisting parity. Consider first the well-known global-workspace theory, according to which a mental state is conscious just in case it is “in” the central cognitive module—the global workspace—and so suitably available for broadcast to the rest of the mind/brain (e.g., Baars 1997; Dehaene et al 2006). If this view were true, why would it be the case that perceptual states, but not belief states, must be globally broadcast? Or consider higher-order theories, on which conscious states are those states of which one is suitably aware of being in (e.g., Armstrong 1968; Lau & Rosenthal 2011). Why would it be the case, if a higher-order theory of consciousness were correct, that must we form higher-order awareness of perceptual but not desire states? Perhaps some explanation is in the offing, but it is far from clear what it may be.

The skeptic might, however, pursue the second horn of our dilemma by denying that the examples we have offered involve unconscious mental states, instead maintaining that they involve either conscious mentality only or subpersonal informational processing. There is, for example, striking evidence that people are often able to predict accurately the results of their implicit-associate tests, suggesting that implicit biases may in fact be conscious (e.g., Hahn et al 2014). But such evidence is hardly dispositive—and there remain good reasons to think that such cognition is unconscious (see, e.g., Berger forthcoming).

Alternatively, the skeptic might take issue with the various experimental results on methodological grounds. Specifically, they might argue that there is something akin to the problem of the criterion that these results face, that explains why these cognitive states appear to be unconscious, when in fact they are conscious but merely unreported. In the implicit-bias case, there clearly is such an explanation available: people are unlikely to report on their own conscious biases, given their social unacceptability. But it is far less clear in the case of Type-1 processing, what the mechanism could be that explains why people fail to report on the mental states involved. Indeed, the whole framework that supports the problem of the criterion makes the most sense in the context of perception, where signal detection theory can be straightforwardly applied. It is much harder to see how it can successfully be applied in nonperceptual cases.

The other option for the skeptic is to accept that such cases are genuinely unconscious, but to urge that they are merely subpersonal. Here the skeptic could pursue analogs of the problem of attribution, arguing that these nonperceptual states fail to meet the criterion for personal-level attribution. A case by case analysis would take us too far afield here, but at first blush it would seem that these states play roughly the same functional roles as their conscious counterparts—they are, for

example, action guiding in ways that seem rational (in the case of implicit bias, see, e.g., Mandelbaum 2016).¹⁰ So it is unclear on what grounds we might deny that they are personal-level. Our point is that the burden of proof is on the skeptic to provide such grounds, given that there at least seem to be many coherent and folk-psychological cases of unconscious mentality.

This dilemma for the strong skeptic is not decisive, but without some reply, the natural conclusion is that it remains plausible that perceptual states may occur unconsciously as well.

8. Conclusions

As we have seen, P&L offer fascinating experimental evidence suggesting that studies of masking in healthy participants wherein discrimination is above chance involve perception that is conscious. Phillips, by contrast, argues that studies wherein discrimination is at chance fail to demonstrate that genuine perception has occurred. Neither criticism is dispositive. We have not shown that any particular study demonstrates unconscious perception, but rather that the reasons to be doubtful are not as strong as they may have initially appeared.

Of course, more convincing evidence may be produced. For example, one could generate a masked-priming experiment wherein the relevant perceptual states figure more clearly in action. Similarly, one could devise a follow-up to P&L's study wherein the hypothesis that the perception involved is conscious is teased apart from the competing hypothesis that the perception involved is unconscious but causes conscious hunches. We leave such explorations to experimentalists. For now, we conclude that the case for unconscious perception remains very much open.¹¹

References

- Armstrong, D. M. 1968. *A Materialist Theory of the Mind*. London: Routledge and Kegan Paul.
- Baars, B. J. 1997. In the theatre of consciousness: Global workspace theory, a rigorous scientific theory of consciousness. *Journal of Consciousness Studies* 4(4): 292-309.
- Bachmann, T. & Francis, G. *Visual Masking: Studying Perception, Attention, and Consciousness*. Oxford: Academic Press.
- Bayne, T., Cleeremans, A., & Wilken, P. (eds.). 2014. *The Oxford Companion to Consciousness*. Oxford: Oxford University Press.

¹⁰ It's worth noting that there are some in the implicit bias literature who deny that implicit biases are full-fledged propositional attitudes such as belief (e.g., Levy 2015), but they do not go so far as to deny that they are personal-level mental states.

¹¹ We thank Hakwan Lau, Marjan Persuh, Ian Phillips, Katia Samoilova, and audiences at the 2107 Intermountain Philosophy Conference and 2018 Pacific Division Meeting of the American Philosophical Association for their helpful comments on this material

- Berger, J. 2014. Mental states, conscious and nonconscious. *Philosophy Compass* 9(6): 392-401.
- Berger, J. 2015. The sensory content of perceptual experience. *Pacific Philosophical Quarterly* 96(4): 446-468.
- Berger, J. forthcoming. Implicit attitudes and awareness. *Synthese* doi: 10.1007/s11229-018-1754-3.
- Block, N. 2016. The Anna Karenina principle and skepticism about unconscious perception. *Philosophy and Phenomenological Research* 93(2): 452–459.
- Burge, T. 2010. *Origins of Objectivity*. Oxford: Oxford University Press.
- Byrne, A. 2001. Intentionalism defended. *Philosophical Review* 110(2): 199-240.
- D'Aloisio-Montilla, N. 2018. Phillips on unconscious perception and overflow. *Philosophia* <https://doi.org/10.1007/s11406-018-0021-7>.
- Danckert, J. and Rossetti, Y. 2005. Blindsight in action: what can the different sub-types of blindsight tell us about the control of visually guided actions? *Neuroscience and Biobehavioral Reviews* 29, 1035-1046.
- Dehaene, S., J.P. Changeux, L. Naccache, J. Sackura, & C. Sergenta. 2006. Conscious, preconscious, and subliminal processing: a testable taxonomy. *Trends in Cognitive Science* 10: 204-11.
- Dennett, D. 1969. *Content and Consciousness*. New York: Routledge and Kegan Paul.
- Drayson, Z. 2012. The uses and abuses of the personal/subpersonal distinction. *Philosophical Perspectives* 26(1): 1-18.
- Fridland, E. 2017. Skill and motor control: intelligence all the way down. *Philosophical Studies* 174(6): 1539-1560.
- Gawronski, B., & Bodenhausen, G. V. 2011. The associative-propositional evaluation model: Theory, evidence, and open questions. *Advances in Experimental Social Psychology* 44: 59-127.
- Greenwald, A. G., McGhee, D. E., & Schwartz, J. K. L. 1998. Measuring individual differences in implicit cognition: The implicit association test. *Journal of Personality and Social Psychology* 74: 1464-1480.
- Hahn, A., Judd, C. M., Hirsh, H. K., & Blair, I. V. 2014. Awareness of implicit attitudes. *Journal of Experimental Psychology: General* 143(3): 1369-1392.
- Knotts, J. D., Lau, H., & Peters, M. A. K. 2018 Continuous flash suppression and monocular pattern masking impact subjective awareness similarly. *Attention, Perception, & Psychophysics* 80(8): 1974-1987.
- Koriat, A. 200. The feeling of knowing: Some metatheoretical implications for consciousness and control. *Consciousness and Cognition* 9: 149-171.
- Lau, H. C., 2009. Volition and the function of consciousness. In N. Murphy, G. F. R. Ellis, T. O'Connor (eds.), *Downward Causation and the Neurobiology of Free Will*. Berlin: Springer
- Lau, H. and Passingham, R.E. (2007) Unconscious activation of the cognitive control system in the human prefrontal cortex. *Journal of Neuroscience* 27: 5805–5811.
- Lau, H. C. & Rosenthal, D. M. 2011. Empirical support for higher-order theories of conscious awareness. *Trends in Cognitive Sciences* 15(8): 365-373.
- Liu, D. & Todorov, E. 2007. Evidence for the flexible sensorimotor strategies predicted by optimal feedback control. *The Journal of Neuroscience* 27(35), 9354-9368.
- Levy, N. 2011. Resisting 'weakness of will'. *Philosophy and Phenomenological Research* 82(1): 134-155.
- Levy, N. 2015. Neither fish nor fowl: Implicit attitudes as patchy endorsements. *Noûs* 49(4): 800–823.
- Mandelbaum, E. 2016. Attitude, inference, association: On the propositional structure of implicit bias. *Noûs* 50(3): 629-658.
- Marcel, A. J. 1983. Conscious and unconscious perception: Experiments on visual masking and

- word recognition. *Cognitive Psychology* 15: 197-237.
- Milner, D., & Goodale, M. 2008. Two visual systems re-viewed. *Neuropsychologia* 46: 774-785.
- Nagel, T. 1974. What is it like to be a bat? *Philosophical Review* 83(4): 435-450.
- Overgaard, M., Feh, K., Mouridsen, K., Bergholt, B., & Cleeremans, A. 2008. Seeing without seeing? Degraded conscious vision in a blindsight patient. *PLoS ONE* 3(8): e3028.
- Persaud, N. Davidson, M., Maniscalco, B., Mobbs, D., Passingham, R. E., Cowey, A., Lau, H. 2011. Awareness-related activity in prefrontal and parietal cortices in blindsight reflects more than superior visual performance. *Neuroimage* 58(2): 605-611.
- Persuh, M., Emmanouil, T. A., & Ro, T. 2016. Perceptual overloading reveals illusory contour perception without awareness of the inducers. *Attention, Perception & Psychophysics* 78(6): 1692-701.
- Persuh, M. 2018. Measuring perceptual consciousness. *Frontiers in Psychology* 8: 2320.
- Peters, M. A., & Lau, H. 2015. Human observers have optimal introspective access to perceptual processes even for visually masked stimuli. *eLife* 4: e09651.
- Peters, M. A., Ro, T., & Lau, H. 2016. Who's afraid of response bias? *Neuroscience of Consciousness* 2016(1).
- Peters, M. A., Kentridge, R., Phillips, I., & Block, N. 2017. Does unconscious perception really exist? Continuing the ASSC20 debate. *Neuroscience of Consciousness* 3(1): nix015.
- Phillips, I. 2011. Perception and iconic memory: What Sperling doesn't show. *Mind & Language* 26(2): 381-411.
- Phillips, I. 2014. Naïve realism and the science of (some) illusions. *Philosophical Topics* 44(2): 353-380.
- Phillips, I. 2016. Consciousness and criterion: On Block's case for unconscious seeing. *Philosophy and Phenomenological Research* 93(2): 419-451.
- Phillips, I. 2018. Unconscious perception reconsidered. *Analytic Philosophy* 59(4): 471-514.
- Phillips, I., & Block, N. 2016. Debate on unconscious perception. In B. Nanay (Ed.), *Current Controversies in Philosophy of Perception* (pp. 165–192). New York: Routledge.
- Pierce, C. A., & Jastrow, J. 1884. On small differences in sensation. *Memoirs of the National Academy of Science* 3: 73-83.
- Pisella, L., Gréa, H., Tilikete, C., Vighetto, A., Desmurget, M., Rode, G., Boisson, D., & Rossetti, Y. 2000. An 'automatic pilot' for the hand in human posterior parietal cortex: toward reinterpreting optic ataxia. *Nature Neuroscience* 3(7), 729-736.
- Qian, M. K., Quinn, P. C., Heyman, G. D., Pascalis, O., Fu, G., & Lee, K. 2017. Perceptual individuation training (but not mere exposure) reduces implicit racial bias in preschool children. *Developmental Psychology* 53(5): 845-859.
- Quilty-Dunn, J. 2016. Iconicity and the format of perception. *Journal of Consciousness Studies* 23(3-4): 255-263.
- Quilty-Dunn, J. 2018. Unconscious perception and phenomenal coherence. *Analysis* <http://doi.org/10.1093/analys/any022>
- Reingold, E. M., & Merikle, P. M. 1988. Using direct and indirect measures to study perception without awareness. *Perception & Psychophysics* 44:563-575.
- Reuter, K. & Sytsma, J. 2018. Unfelt pain. *Synthese* <https://doi.org/10.1007/s11229-018-1770-3>
- Rosenthal, D. M. 1990. On being accessible to consciousness. *Behavioral and Brain Sciences* 13(4): 621-622.
- Rosenthal, D. M. 2012. Higher-order awareness, misrepresentation and function. *Philosophical Transactions of the Royal Society of London B: Biological Sciences* 367(1594): 1424-1438.
- Rosenthal, D. M. forthcoming. Consciousness and confidence. *Neuropsychologia* <https://doi.org/10.1016/j.neuropsychologia.2018.01.018>.
- Seitz, A.R., Kim, D., and Watanabe, T. 2009. Rewards evoke learning of unconsciously processed

- visual stimuli in adult humans. *Neuron* 61(5): 700–707.
- Shanks, D. R. 2017. Regressive research: The pitfalls of post hoc data selection in the study of unconscious mental processes. *Psychonomic Bulletin & Review* 23(3): 752-775.
- Soto, D., Mantyla, T., & Silvanto, J. 2011. Working memory without consciousness. *Current Biology* 21(22): R912-913.
- Sperling, G. 1960. The information available in brief visual presentations. *Psychological Monographs: General and Applied* 74(11):1– 29.
- Stein, T., Kaiser, D., & Hesselmann, G. 2016. Can working memory be non-conscious? *Neuroscience of Consciousness* 2016(1): 1-3.
- Stanislaw, H. & Todorov, N. 1999. Calculation of signal detection theory measures. *Behavior Research Methods, Instruments, & Computers: A Journal of the Psychonomic Society* 31: 137-49.
- Tononi, G., M. Boly, M. Massimini, & C. Koch. 2016. Integrated information theory: From consciousness to its physical substrate. *Nature Reviews Neuroscience* 17(7): 450-461.
- Williams L. M., Das, P., Lidell, B. J., Kemp, A. H., Rennie, C. J., Gordon, E. 2006. Mode of functional connectivity in amygdala pathways dissociates level of awareness for signals of fear. *Journal of Neuroscience* 26(36): 9264-9271.
- Weiskrantz L. 1986. *Blindsight: A Case Study and Implications*. Oxford: Oxford University Press.
- Weiskrantz, L. 2009. Is blindsight just degraded normal vision? *Experimental Brain Research* 192(3): 413-416.